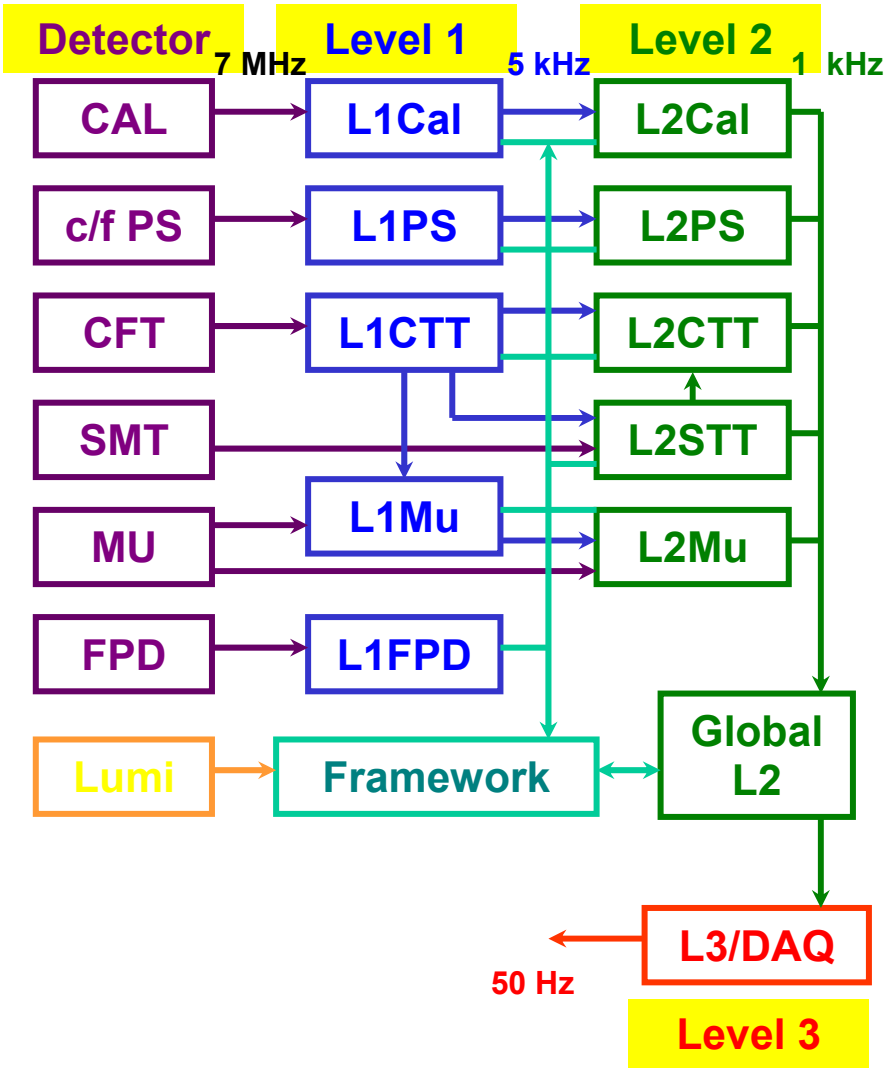

TRIGGER UPGRADE

(This page intentionally left blank)

TRIGGER UPGRADE CONTENTS

1 Introduction

A powerful and flexible trigger is the cornerstone of a modern hadron collider experiment. It dictates what physics processes can be studied properly and what is ultimately left unexplored. The trigger must offer sufficient flexibility to respond to changing physics goals and new ideas. It should allow the pursuit of complementary approaches to a particular event topology in order to maximize trigger efficiency and allow measurement of trigger turn-on curves. Adequate bandwidth for calibration, monitoring, and background samples must be provided in order to calibrate the detector and control systematic errors. If the trigger is not able to achieve sufficient selectivity to meet these requirements, the capabilities of the experiment will be seriously compromised.

The DØ Run IIb Trigger Upgrade is designed to meet these goals within the context of the physics program described in Part II and the challenges of triggering at the high instantaneous luminosities that will be present in Run IIb. The upgrade will allow DØ to select with high efficiency the wide variety of data samples required for the Higgs search and the high-p_T physics program, while providing sufficient background rejection to meet constraints imposed by the readout electronics and DAQ system.

The DØ Run IIb Trigger Upgrade is designed for operation at a peak luminosity of 5×10^{32} $cm^{-2}s^{-1}$ with 132 ns bunch spacing. We have also investigated operating at a luminosity of 2×10^{32} cm^{-2} sec^{-1} with 396 ns bunch spacing. In this operating mode, luminosity leveling is used to hold the luminosity at 2×10^{32} cm^{-2} sec^{-1} and the achievable integrated luminosity is expected to be the same as if there were no leveling and an initial luminosity of about 3.4×10^{32} cm^{-2} sec^{-1}. Both modes of operation yield an average of 5 minimum bias interactions accompanying the high-p_T interaction and require a factor of ~2.5 increase in trigger rejection over the Run IIa design.

Laboratory guidance for Run IIb is that a luminosity of 2×10^{32} cm^{-2} sec^{-1} with 396 ns bunch spacing and luminosity leveling is the baseline plan, but that CDF and DØ should have the capability of operating at higher instantaneous luminosities with either 132 ns or 396 ns bunch spacing should luminosity leveling not meet expectations. Our proposed trigger upgrade is consistent with this guidance.

We will retain the present trigger architecture with three trigger levels. The Level 1 (L1) trigger employs fast, deterministic algorithms, generating an accept/reject decision every 132 ns. Each detector element (except for the Silicon Tracker) has a dedicated trigger processor that implements the trigger algorithm and supplies “And/Or” input bits to the Trigger Framework, which makes the L1 trigger decision. The Level 2 (L2) trigger utilizes both DSP and general-purpose processors to implement more complex algorithms with variable processing time and correlate information from different detectors into the trigger decision. An important part of the L2 trigger is the Silicon Track Trigger (STT), which performs track fits using fiber and silicon tracker hits and allows triggering on displaced vertices. The Level 3 (L3) trigger is based on a farm of high-

performance commodity processors, each analyzing a complete event. While the L1 and L2 trigger rely on dedicated trigger data paths, the L3 trigger utilizes the DAQ readout to collect all event data in a L3 processing node. Since the L3 trigger is an integral part of the DAQ, we defer discussion of the L3 trigger upgrade to Part 4: DAQ/Online Computing of the TDR.

We cannot accommodate the higher Run IIb luminosity by simply increasing trigger rates. The L1 trigger rate is limited to a peak rate of ~5 kHz by readout deadtime. The L2 trigger rate is limited to a peak rate of ~1 kHz by the calorimeter digitization time. Finally, we have set a goal of ~50 Hz for the average L3 trigger rate to limit the strain on (and cost of) data storage and offline computing.

The above L1 and L2 rate limits remain the same in Run IIb as in Run IIa. Thus, for Run IIb we must increase the L1 trigger rejection by a factor of 2.5 and maintain the current L2 rejection factor of 5. Since Run IIb will focus primarily on high-p_T physics processes, we expect some bandwidth will be freed by reducing the trigger rate devoted to low-p_T processes. As we show in Section 2, this reduction is not sufficient to meet our rate limitations, nor does it address the difficulties in triggering efficiently on some important high-p_T processes. Only by upgrading the trigger will we have a reasonable level of confidence in our ability to acquire the data samples needed to carry out the Run IIb physics program.

Potential Run IIb trigger upgrades are further limited by the relatively short time available. The upgrade must be ready for installation in 2005 so that installation of the trigger and silicon upgrades can proceed in parallel. Thus, we have been careful to limit the number and scope of the proposed Run IIb trigger upgrades to those that are necessary and sufficient to meet the Run IIb goals.

In the sections below, we describe the technical design of the Run IIb trigger upgrade. Section 2 provides an overview of the trigger architecture and some of the triggering challenges that must be overcome for Run IIb. Section 3 describes the design of the L1 track trigger, which generates track-based triggers and provides tracking information to several other trigger systems. Section 4 describes the design of a new L1 calorimeter trigger that will replace the current trigger (one of the few remaining pieces of Run 1 electronics in DØ). The calorimeter upgrade will employ digital filtering to associate energy with the correct beam crossing in the Run IIb environment and provide the capability of clustering energy from multiple trigger towers. It will also allow improved e/γ/τ triggers that make fuller use of the calorimeter (HAD/EM, cluster shape/size, isolation) and tracking information. Section 5 describes the calorimeter-track matching system, which is based on the existing muon-track matching system. These improvements to the L1 trigger will significantly reduce the rate for multijet background by sharpening trigger thresholds and improving particle identification. Section 6 describes the upgrade of the L2 βeta processors to provide additional computational power at L2. Section 7 describes the changes to the L2 Silicon Track Trigger needed to accommodate the new silicon tracker being built for Run IIb. Lastly, Section 8 summarizes the trigger part of the Technical Design Report.

2 Triggers, Trigger Terms, and Trigger Rates

The primary feature driving the design of the Run IIb trigger elements is the higher rates associated with the increased instantaneous luminosity that will be delivered to the experiments. The rate of events accepted by the Level 1 trigger still must be limited to ~5 kHz to maintain acceptable deadtime levels, so the overall aim of the Level 1 upgrade is to increase the rejection by a factor of at least 2.5.

At 2 TeV, the inelastic proton-antiproton cross section is very large, about 50 mb. At Run 2 luminosities, this results in interaction rates of ~25 MHz, with multiple interactions occurring in most beam crossings. Virtually all of these events are without interest to the physics program. In contrast, at these luminosities W bosons are produced at a few Hz and a few top quark pairs are produced per hour. It is evident that sophisticated triggers are necessary to separate out the rare events of physics interest from the overwhelming backgrounds. Rejection factors of nearly 10^6 must be achieved in decision times of a few milliseconds.

The salient features of interesting physics events naturally break down into specific signatures which can be sought after in a programmable trigger. The appearance in an event of a high p_T lepton, for example, can signal the presence of a W or a Z. Combined with jets containing *b* quark tags, the same lepton signature could now be indicative of top quark pair production or the Higgs. Leptons combined instead with missing energy is a classic SUSY discovery topology, etc. The physics "menu" of Run 2 is built on the menu of signatures and topologies available to the trigger. In order for the physics program to succeed, these fundamental objects must remain un-compromised at the highest luminosities. The following paragraphs give a brief overview of the trigger system and a sampling of the physics impact of the various combinations of trigger objects.

2.1 Overview of the DØ Run IIa Trigger System

The DØ trigger system for Run II is divided into three levels of increasing complexity and capability. The Level 1 (L1) trigger is entirely implemented in hardware (see Figure 1). It looks for patterns of hits or energy deposition consistent with the passage of high energy particles through the detector. The calorimeter trigger tests for energy in calorimeter towers above pre-programmed thresholds. Hit patterns in the muon system and the Central Fiber Tracker (CFT) are examined to see if they are consistent with charged tracks above various transverse momentum thresholds. These tests take up to 3.5 μs to complete, the equivalent of 27 beam crossings. Since ~10 μs of deadtime for readout is incurred following a L1 trigger, we have set a maximum L1 trigger rate of 5 kHz.

Each L1 system prepares a set of terms representing specific conditions that are satisfied (e.g. 2 or more CFT tracks with p_T above 3 GeV). These hardware terms are sent to the L1 Trigger Framework, where specific triggers are formed from combinations of terms (e.g. 2 or more CFT tracks with p_T above 3 GeV AND

2 or more EM calorimeter clusters with energy above 10 GeV). Using firmware, the trigger framework can also form more complex combinations of terms involving ORs of hardware terms (e.g. a match of preshower and calorimeter clusters in any of 4 azimuthal quadrants). The Trigger Framework has capacity for 256 hardware terms and about 40 firmware terms.

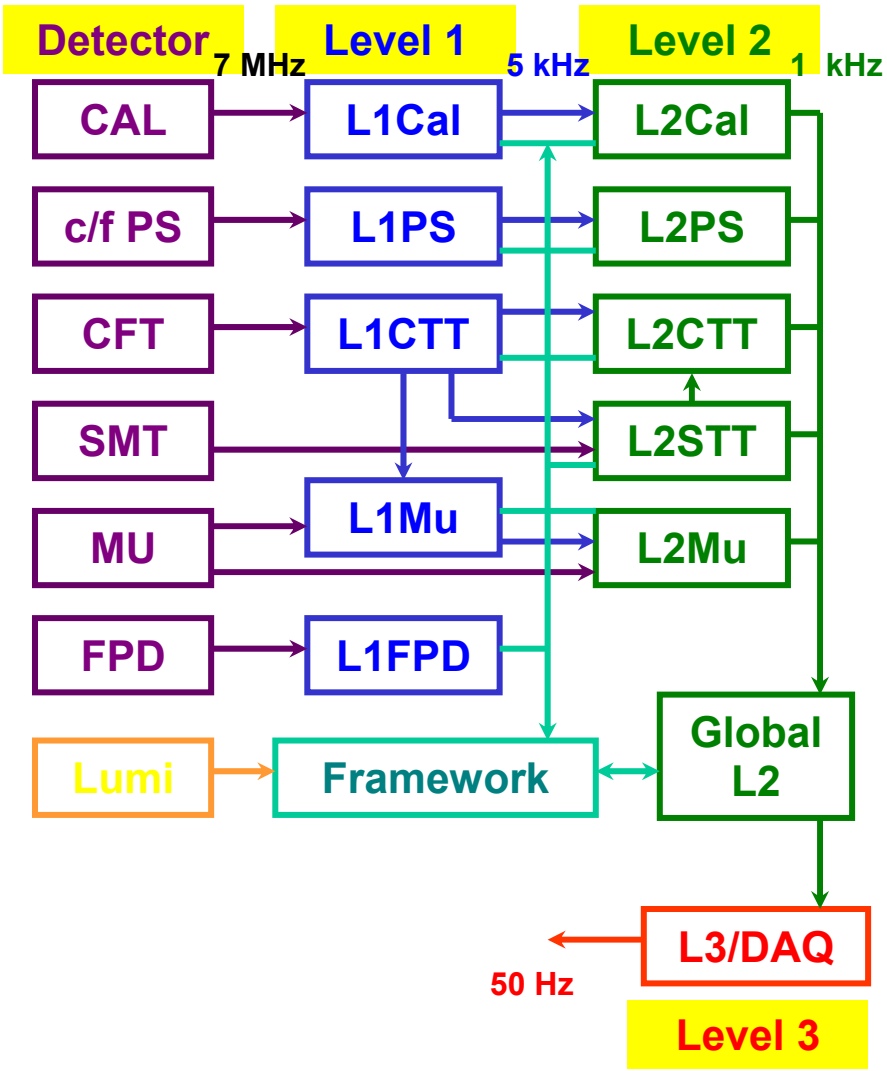

Figure 1. Block diagram of the trigger system, indicating the individual trigger processors that comprise each level.

The Level 2 trigger (L2) takes advantage of the spatial correlations and more precise detector information to further reduce the trigger rate. The L2 system consists of dedicated preprocessors, each of which reduces the data from one detector subsystem (calorimeter, muon, CFT, preshowers, and SMT). A global L2 processor takes the individual elements and assembles them into physics "objects" such as muons, electrons, or jets. The Silicon Track Trigger (STT) introduces the precise track information from the SMT to look for large impact parameter tracks from *b*-quark decays. Some pipelining is necessary at L2 to

meet the constraints of the 100 μs decision time. L2 can accept events and pass them on to Level 3 at a rate of up to 1 kHz.

The Level 3 (L3) trigger consists of a farm of fast, high-level computers (PCs) which perform a simplified reconstruction of the entire event. Even within the tight time budget of 25 ms, this event reconstruction will allow the application of algorithms in the trigger with sophistication very close to that of the offline analyses. Events that satisfy desired characteristics will then be written out to a permanent storage medium. The average L3 output for Run IIa is 50 Hz and is largely dictated by downstream computing limits.

2.2 Leptonic Triggers

Leptons provide the primary means of selecting events containing W and Z bosons. They can also tag *b* quarks through their semileptonic decays, complementing the more efficient (but only available at Level 2 through the STT) lifetime selection. The impact of the purely leptonic tag is seen most strongly in the measurements of the W mass, the W and Z production cross sections, and the W width, since the events containing W and Z bosons are selected solely by requiring energetic leptons. The increased statistics provided by Run IIb should allow for a significant improvement in the precision of the precision of these measurements, complementing the direct searches in placing more stringent constraints on the Standard Model.

In addition to their inherent physics interest, leptonic signals will play an increasingly important role in the calibration of the energy and momentum scales of the detectors, which is crucial for the top quark and W mass measurements. This will be accomplished using $Z \rightarrow e^+e^-$, $\Upsilon \rightarrow e^+e^-$, and $J/\Psi \rightarrow e^+e^-$ for the electromagnetic calorimeter energy scale and the corresponding muon decays for the momentum scale. Since the trigger bandwidth available for acquiring calibration samples must be non-zero, another set of constraints is imposed on the overall allocation of trigger resources.

2.3 Leptons plus Jets

During Run I, lepton-tagged decays of the W bosons and *b* quarks played an essential role in the discovery of the top quark and were exploited in the measurements of the top mass and production cross section. The new capability provided by the STT to tag *b* quark decays on-line will allow the collection of many thousands of $t\bar{t}$ pairs in the channel $t\bar{t} \rightarrow \ell\nu$+jets with one *b*-tagged jet. This will be sufficient to allow the study of top production dynamics as well as the measurement of the top decay branching fractions. The precision in measuring the top quark mass will ultimately be limited by our ability to control systematic errors, and the increase in statistics for Run IIb will allow the reduction of several key systematic errors for this channel as well as for the channel $t\bar{t} \rightarrow \ell\nu\ell'\nu$+jets. One of these, the uncertainty in the jet energy scale, can be reduced by understanding the systematics of the direct reconstruction of W or Z boson decays into jets. The most promising channel in this case is the decay $Z \rightarrow b\bar{b}$,

in which secondary vertex triggers can provide the needed rejection against the dominant two-jet background.

2.4 Leptons/Jets plus Missing E_T

Events containing multiple leptons and missing energy are often referred to as the "gold-plated" SUSY discovery mode. These signatures, such as three leptons plus missing energy, were explored in Run I to yield some of the most stringent limits on physics beyond the Standard Model. These investigations will be an integral part of the search for new physics in Run 2. Missing energy is characteristic of any physics process where an invisible particle, such as an energetic neutrino or a massive stable neutral particle, carries away a large fraction of the available energy. Missing energy combined with leptons/photons or jets can be a manifestation of the presence of large extra dimensions, different SUSY configurations, or other new physics beyond the Standard Model.

2.5 Triggers for Higgs Searches

One of the primary goals of the Run IIb physics program will be to exploit the delivered luminosity as fully as possible in search of the Higgs boson up to the highest accessible Higgs mass[1]. Since even a delivered luminosity of 15 fb^{-1} per experiment may not lead to a statistically significant discovery, the emphasis will be on the combination of as many decay channels and production mechanisms as possible to maximize the prospects for Higgs discovery. For the trigger, this implies that flexibility, ease of monitoring, and selectivity will be critical issues.

Coverage of the potential window of discovery is provided by the decay channel $H \rightarrow b\bar{b}$ at low masses, and by $H \rightarrow W^{(*)}W$ at higher masses. In the first case, the production mechanism with the highest sensitivity will probably be in the mode $p\bar{p} \rightarrow$ WH. For leptonic W decays, the leptons can be used to trigger on the events directly. If the W decays hadronically, however, the four jets from the $q\bar{q}b\bar{b}$ final state will have to be pulled out from the large QCD backgrounds. Tagging *b* jets on-line will provide a means to select these events and ensure that they are recorded. Of course, three or four jets with sufficient transverse energy are also required. Another decay mode with good sensitivity is $p\bar{p} \rightarrow$ ZH, where the Z decays to leptons, neutrinos, or hadrons. From a trigger perspective, the case where the Z decays hadronically is identical to the WH all-hadronic final state. The final state $ZH \rightarrow \nu\bar{\nu}b\bar{b}$, however, provides a stringent test for the jet and missing E_T triggers, since the final state is only characterized by two modest *b* jets and missing energy.

Recently, the secondary decay mode $H \rightarrow \tau^+\tau^-$ has come under scrutiny as a means of bolstering the statistics for Higgs discovery in the low mass region. A trigger that is capable of selecting hadronic tau decays by means of isolated, stiff tracks or very narrow jets will give access to the gluon-fusion production mode gg

[1] Report of the Higgs Working Group of the Tevatron Run 2 SUSY/Higgs Workshop, M. Carena *et al*, hep-ph/0010338.

$\rightarrow H \rightarrow \tau^+ \tau^-$ for lower Higgs masses. A preliminary analysis has demonstrated[2] that the inclusion of this mode could reduce by 35% the luminosity required for a discovery/exclusion of the SM Higgs boson up to masses of 140 GeV, making it clearly worth pursuing. This mode can also be important in some of the large tanβ SUSY scenarios, where the Higgs coupling to $b\bar{b}$ is reduced, leaving $H \rightarrow \tau^+ \tau^-$ as the dominant decay mode for the lightest Higgs.

The higher Higgs mass regime will be covered by selecting events from $p\bar{p} \rightarrow H \rightarrow W^{(*)}W$ with one or two high-energy leptons from the $W \rightarrow \ell\nu$ decay. This decay mode thus requires a trigger on missing E_T in addition to leptons or leptons plus jets. Supersymmetric Higgs searches will require triggering on final states containing 4 b-quark jets. This will require jet triggers at L1 followed by use of the STT to select jets at L2.

2.6 Trigger Menu and Rates

As even this cursory review makes clear, the high-p_T physics menu for Run IIb requires efficient triggers for jets, leptons (including taus, if possible), and missing E_T at Level 1. The STT will be crucial in selecting events containing *b* quark decays; however, its rejection power is not available until Level 2, making it all the more critical that the Level 1 system be efficient enough to accept all the events of interest without overwhelming levels of backgrounds.

In an attempt to set forth a trigger strategy that meets the physics needs of the experiment, the Run 2 Trigger Panel suggested a preliminary set of Trigger Terms for Level 1 and Level 2 triggers[3]. In order to study the expected rates in Run IIb, we have simulated an essential core of triggers which cover the essential high-pt physics signatures: Higgs boson produced in association with W and Z bosons with Higgs decays to b-bbar, Higgs production and decay to tau leptons, top quark decays in leptonic and semi-leptonic channels, inclusive W and Z boson decays into lepton and muons. The simple triggers we have currently implemented at Level 1 for Run IIa will not be able to cope with the much higher occupancies expected in Run IIb without a drastic reduction in the physics scope of the experiment and/or prescaling of important physics triggers. Our rate studies have used QCD jets samples in order to determine the effects of background, including multiple low-p_T minimum bias events superimposed on the dominant processes. Table 1 shows the effect of the Run IIb trigger upgrades described below on this selection of L1 triggers for a luminosity of 2×10^{32} cm^{-2} sec^{-1} with 396 ns bunch spacing. Without the upgrade, the L1 trigger rate will far exceed the allowed bandwidth of 5 kHz. With the upgrade, the L1 trigger rate meets the bandwidth constraint with some headroom for the additional triggers that will be required for a robust Run IIb physics program.

---

[2] A. Belyaev, T. Han, and R. Rosenfeld *"gg→H→ττ at the Upgraded Fermilab Tevatron"*, hep-ph/0204210, April 2002.

[3] The report of the Run 2 Trigger Panel can be found at http://d0server1.fnal.gov/projects/run2b/Trigger/Docs/Trigger_Panel_Report.ps.

Table 1. Trigger rates for an example trigger menu representing a full spectrum of Run 2 physics channels. Representative physics channels for each of the triggers is indicated. The rates for each of these triggers with the design Run IIa trigger and the Run IIb upgraded trigger are also shown.

Trigger	Example Physics Channels	L1 Rate (kHz) (no upgrade)	L1 Rate (kHz) (with upgrade)
EM (1 EM TT > 10 GeV)	$W \to e\nu$ $WH \to e\nu jj$	1.3	0.7
Di-EM (1 EM TT > 7 GeV, 2 EM TT > 5 GeV)	$Z \to ee$ $ZH \to eejj$	0.5	0.1
Muon (muon p_T > 11 GeV + CFT Track)	$W \to \mu\nu$ $WH \to \mu\nu jj$	6	0.4
Di-Muons (2 muons p_T > 3 GeV + CFT Tracks)	$Z \to \mu\mu$, $J/\psi \to \mu\mu$ $ZH \to \mu\mu jj$	0.4	< 0.1
Electron + Jets (1 EM TT > 7 GeV, 2 Had TT > 5 GeV)	$WH \to e\nu + jets$ $t\bar{t} \to e\nu + jets$	0.8	0.2
Muon + Jet (muon p_T > 3 GeV, 1 Had TT > 5 GeV)	$WH \to \mu\nu + jets$ $t\bar{t} \to \mu\nu + jets$	< 0.1	< 0.1
Jet+MET (2 TT > 5 GeV, Missing E_T> 10 GeV)	$ZH \to \nu\bar{\nu}b\bar{b}$	2.1	0.8
Muon + EM (muons p_T > 3 GeV + CFT track, 1 EM TT > 5 GeV)	$H \to WW$ $H \to ZZ$	< 0.1	< 0.1
Single Isolated Track (1 Isolated CFT track, p_T > 10 GeV)	$H \to \tau\tau$ $W \to \mu\nu$	17	1.0
Di-Track (1 isolated tracks p_T > 10 GeV, 2 tracks p_T > 5 GeV, 1 track with EM energy)	$H \to \tau\tau$	0.6	<0.1
Total Rate		**28**	**3.6**

We now turn to describing the upgrades to the trigger system that will enable us to cope with the large luminosities and high occupancies of Run IIb.

3 Level 1 Tracking Trigger

The Level 1 Central Tracking Trigger (CTT) plays a crucial role in the full range of L1 triggers. In this section, we outline the goals for the CTT, provide an overview of the performance and implementation of the present track trigger, and describe the proposed Run IIb CTT upgrade.

3.1 Goals

The goals for the CTT include providing stand-alone track triggers, combining tracking and preshower information to identify electron and photon candidates, and generating track lists that allow other trigger systems to perform track matching. The latter is a critical part of the L1 muon trigger. We briefly discuss these goals below.

3.1.1 Track Triggers

The CTT provides various Level 1 trigger terms based on counting the number of tracks whose transverse momentum (p_T) exceeds a threshold. Track candidates are identified in the axial view of the Central Fiber Tracker (CFT) by looking for predetermined patterns of hits in all 8 fiber doublet layers. Four different sets of roads are defined, corresponding to p_T thresholds of 1.5, 3, 5, and 10 GeV, and the number of tracks above each threshold can be used in the trigger decision. For example, a trigger on two high p_T tracks could require two tracks with p_T>5 GeV and one track with p_T>10 GeV.

Triggering on isolated tracks provides a complementary approach to identifying high-p_T electron and muon candidates, and is potentially useful for triggering on hadronic tau decays. To identify isolated tracks, the CTT looks for additional tracks within a 12º region in azimuth (ϕ).

3.1.2 Electron/Photon Identification

Electron and photon identification is augmented by requiring a significant energy deposit in the preshower detector. The Central Preshower (CPS) and Forward Preshower (FPS) detectors utilize the same readout and trigger electronics as the fiber tracker, and are included in the discussion of tracking triggers. Clusters found in the axial layer of the CPS are matched in phi with track candidates to identify central electron and photon candidates. The FPS cannot be matched with tracks, but comparing energy deposits before/after the lead radiator allows photon and electron candidates to be distinguished.

3.1.3 Track Matching

Track candidates found in the CTT are important as input to several other trigger systems. CTT information is used to correlate tracks with other detector measurements and to serve as seeds for pattern recognition algorithms.

The Level 1 muon trigger matches CTT tracks with hits in the muon detector. To meet timing requirements, the CTT tracks must arrive at the muon trigger on the same time scale as the muon proportional drift tube (PDT) information becomes available.

The current Level 1 trigger allows limited azimuthal matching of tracking and calorimeter information at the quadrant level (see Section 2.1). Significantly increasing the flexibility and granularity of the calorimeter track matching is an integral part of the proposed modifications for Run IIb (see Section 5).

The L2 Silicon Track Trigger (STT) uses tracks from the CTT to generate roads for finding tracks in the Silicon Microstrip Tracker (SMT). The precision of the SMT measurements at small radius, combined with the larger radius of the CFT, allows displaced vertex triggers, sharpening of the momentum thresholds for track triggers, and elimination of fake tracks found by the CTT. The momentum spectrum for b-quark decay products extends to low p_T. The CTT therefore aims to provide tracks down to the lowest p_T possible. The Run IIa CTT generates track lists down to $p_T \approx 1.5$ GeV. The CTT tracks must also have good azimuthal (ϕ) resolution to minimize the width of the road used by the STT.

In addition to the track lists sent to the STT, each portion of the L1 track trigger (CFT, axial CPS, and FPS) provides information for the Level 2 trigger decision. The stereo CPS signals are also sent to L2 to allow 3-D matching of calorimeter and CPS signals.

3.2 Description of Current Tracking Trigger

We have limited our consideration of potential track trigger upgrades to those that preserve the overall architecture of the current tracking trigger. The sections below describe the tracking detectors, trigger segmentation, trigger electronics, outputs of the track trigger, and the trigger algorithms that have been developed for Run IIa.

3.2.1 Tracking Detectors

The CFT is made of scintillating fibers mounted on eight low-mass cylinders. Each of these cylinders supports four layers of fibers arranged into two doublet layers. The innermost doublet layer on each cylinder has its fibers oriented parallel to the beam axis. These are referred to as axial doublet layers. The second doublet layer has its fibers oriented at a small angle to the beam axis, with alternating sign of the stereo angle. These are referred to as stereo doublet layers. Only the axial doublet layers are incorporated into the current L1 CTT. Each fiber is connected to a visible light photon counter (VLPC) that converts the light pulse to an electrical signal.

The CPS and FPS detectors are made of scintillator strips with wavelength-shifting fibers threaded through each strip. The CPS has an axial and two stereo layers mounted on the outside of the solenoid. The FPS has two stereo layers in front of a lead radiator and two stereo layers behind the radiator. The CPS/FPS fibers are also read out using VLPCs.

3.2.2 CTT Segmentation

The CTT is divided in ϕ into 80 Trigger Sectors (TS). A single TS is illustrated schematically in Figure 2. To find tracks in a given sector, information is needed from that sector, called the home sector, and from each of its two neighboring

sectors. The TS is sized such that the tracks satisfying the lowest p_T threshold (1.5 GeV) is contained within a single TS and its neighbors. A track is 'anchored' in the outermost (H) layer. The ϕ value assigned to a track is the fiber number at the H layer. The p_T value for a track is expressed as the fiber offset in the innermost (A) layer from a radial straight-line trajectory.

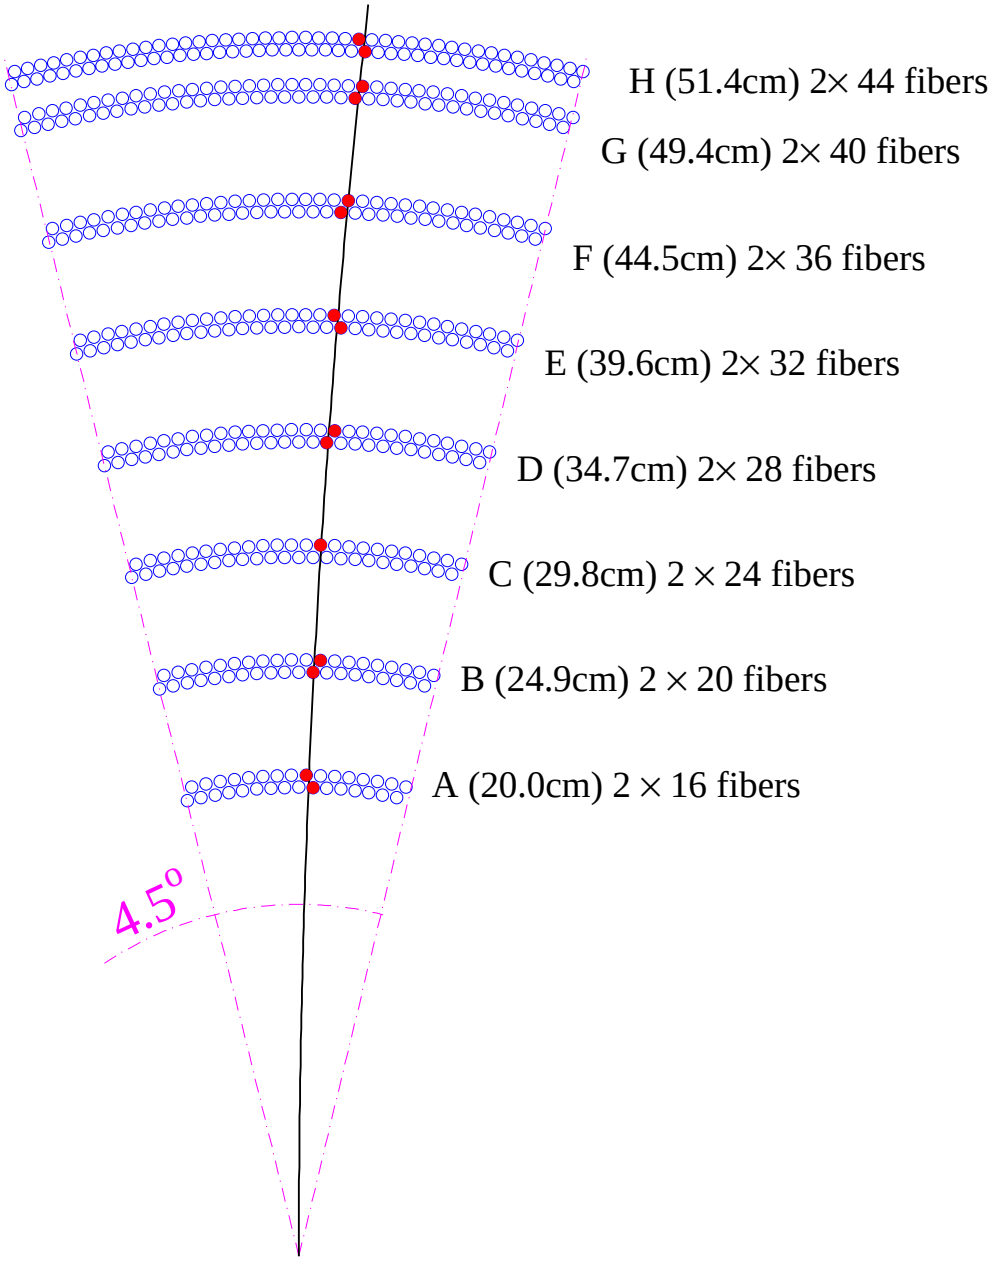

Figure 2. Illustration of a CTT trigger sector and the labels assigned to the eight CFT cylinders. Each of the 80 trigger sectors has a total of 480 axial fibers.

The home sector contains 480 axial fibers. A further 368 axial fibers from each of the neighbors, the 'next' and 'previous', sectors are sent to each home

sector to find all the possible axial tracks above the p_T threshold. In addition, information from 16 axial scintillator strips from the CPS home sector and 8 strips from each neighboring sector are included in the TS for matching tracks and preshower clusters.

3.2.3 Tracking Algorithm

The tracking trigger algorithm currently implemented is based upon hits constructed from pairs of neighboring fibers, referred to as a "doublet". Fibers in doublet layers are arranged on each cylinder as illustrated in Figure 3. In the first stage of the track finding, doublet layer hits are formed from the individual axial fiber hits. The doublet hit is defined by an OR of the signals from adjacent inner and outer layer fibers in conjunction with a veto based upon the information from a neighboring fiber. In Figure 3, information from the first fiber on the left in the upper layer (fiber 2) would be combined by a logical OR with the corresponding information for the second fiber from the left on the lower layer (fiber 3). This combination would form a doublet hit unless the first fiber from the left in the lower layer (fiber 1) was also hit. Without the veto, a hit in both fiber 2 and fiber 1 would result in two doublet hits.

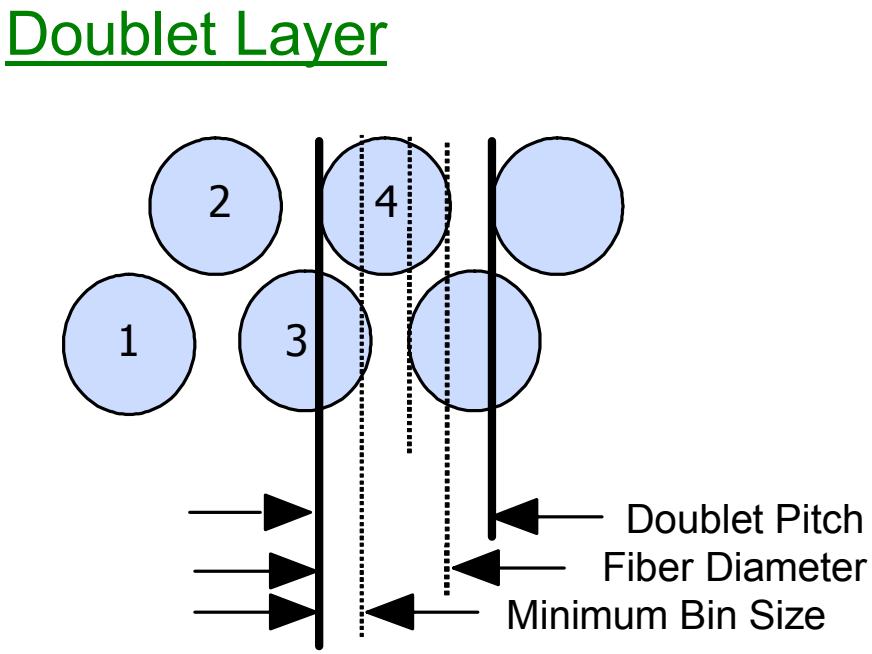

Figure 3. Sketch illustrating the definition of a fiber doublet. The circles represent the active cross sectional areas of individual scintillating fibers. The boundaries of a doublet are shown via the thick black lines. The dotted lines delineate the four distinguishable regions within the doublet.

The track finding within each sector is straightforward. Each pattern of eight doublets that is hit by a track with a given p_T and ϕ is represented by a block of logic. In the following, we will refer to this logic as an "equation". Each equation

forms an 8-fold AND. If all the fibers in the pattern were hit then all 8 inputs of the AND are TRUE and the result is a TRUE. This logic is implemented in large field programmable gate arrays (FPGA). Each TS has 44 ϕ bins corresponding to the 44 H layer doublets in a sector and 20 possible p_T bins for about 12 different routes through the intermediate layers with fixed phi and p_T. This results in about 17K equations per TS. For each TS, up to six tracks with the highest p_T are reported to the trigger.

3.3 Performance with the Run IIa Tracking Trigger

We have simulated the rates to be expected for pure track triggers in Run IIb, taking into account the additional minimum bias events within the beam crossing of interest due to the increased luminosity.

3.3.1 Simulations of the Run IIa trigger

Under Run IIa conditions, the current track trigger performs very well in simulations. For example, for a sample of simulated muons with $p_T > 50$ GeV/c, we find that 97% of the muons are reconstructed correctly; of the remaining 3%, 1.9% of the tracks are not reconstructed at all and 1.1% are reconstructed as two tracks due to detector noise. (As the background in the CFT increases, due to overlaid events, we expect the latter fraction to get progressively higher). Since the data-taking environment during Run IIb will be significantly more challenging, it is important to characterize the anticipated performance of the current trigger under Run IIb conditions.

To test the expected behavior of the current trigger in the Run IIb environment, the existing trigger simulation code was used with an increased number of overlaid minimum bias interactions. The minimum bias interactions used in this study were generated using the ISAJET Monte Carlo model. Based on studies of detector occupancy and charged track multiplicity in minimum-bias events, we expect that this should give a worst-case scenario for the Run IIb trigger.

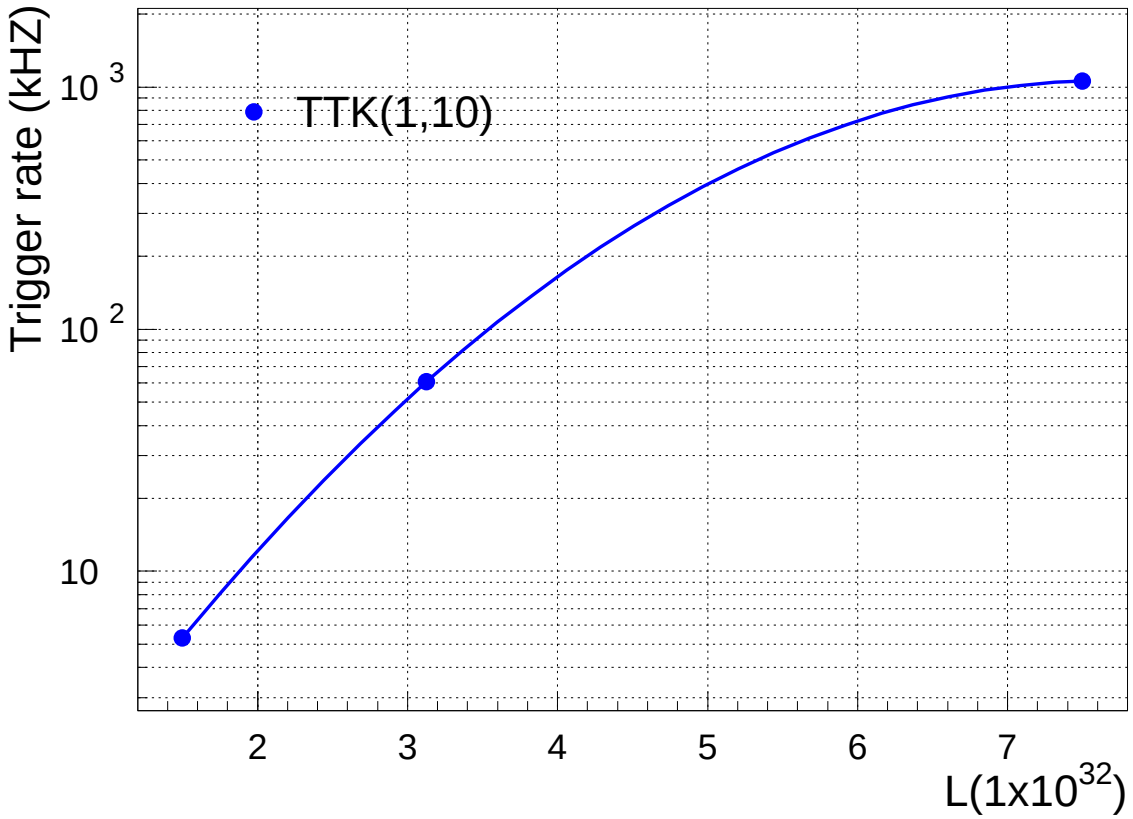

Figure 4. Track trigger rate as a function of the number of underlying minimum bias interactions. TTK(2,10) is a trigger requiring 2 tracks with transverse momentum greater than 10 GeV.

Figure 4 shows the rate for a trigger requiring two tracks with $p_T > 10$ GeV as a function of the number of underlying minimum bias interactions and hence luminosity. During Run IIb, we expect that the mean number of underlying interactions will be about 5. Figure 4 shows that the tracking trigger rate for the current trigger version is expected to rise dramatically due to accidental hit combinations yielding fake tracks. This results in an increasingly compromised tracking trigger.

Figure 5 shows the probability for three specific track trigger terms to be satisfied in a given crossing. They are strongly dependent upon the number of underlying minimum bias interactions. These studies indicate that a track trigger based upon the current hardware will be severely compromised under Run IIb conditions. Not shown in the figure, but even more dramatic, is the performance of the 5 GeV threshold track trigger. This is satisfied in more than 95% of beam crossings with 5 minbias interactions. It will clearly not be possible to run the current stand-alone track trigger in Run IIb. But much worse, the information available to the muon trigger, electron trigger, and STT becomes severely compromised by such a high rate of fake high-p_T tracks.

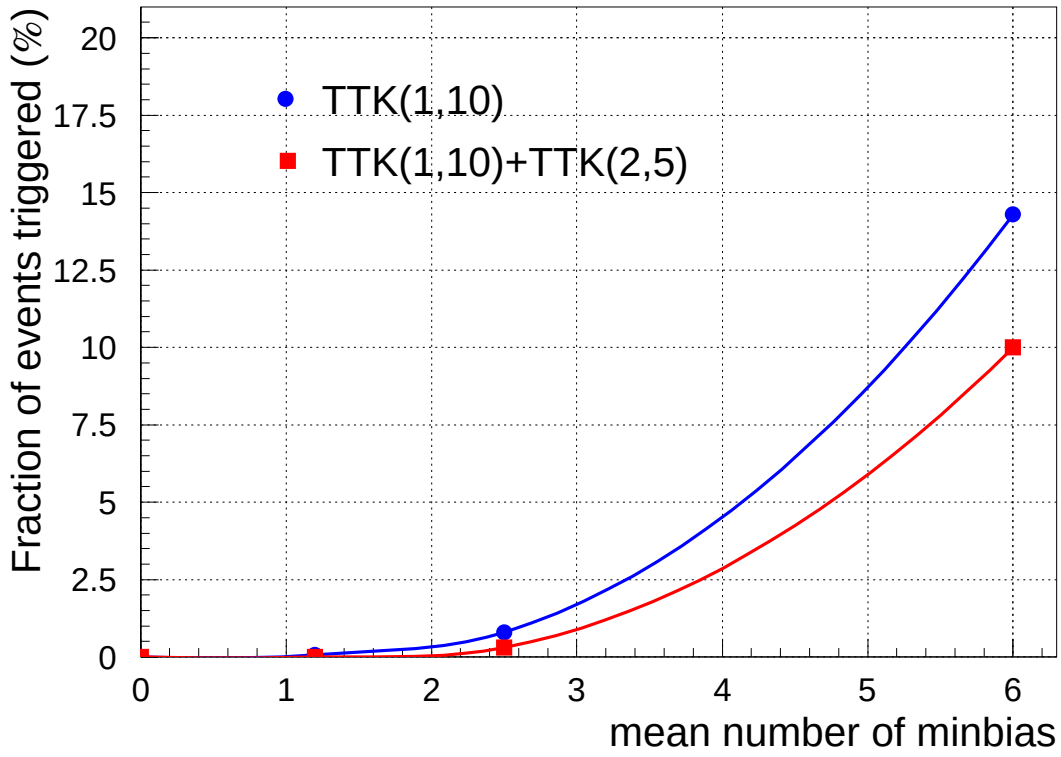

Figure 5. The fraction of events satisfying several track term requirements as a function of the number of minimum bias events overlaid. TTK(n,p_T) is a trigger requiring n tracks with transverse momentum greater than p_T.

Based upon these simulations, we believe that the significant number of multiple interactions in Run IIb and the large CFT occupancy fractions they induce will compromise performance of the current tracking trigger.

3.3.2 Studies of CFT Occupancy

At this time, the Run IIa L1CTT track trigger is still being commissioned. Recently, however, the final CFT read-out electronics began to function at close to their expected signal-to-noise performance. Given this, we can study the CFT occupancy due to noise and minimum bias events. Since we expect minimum bias events to dominate the detector occupancy during high-luminosity running, it is crucial to understand the detector's response to these events.

Our baseline simulation uses the Pythia Monte Carlo generator tuned to CDF Run I data in order to generate minimum bias events. The minimum bias model includes contributions from two-body quark-quark, quark-gluon, and gluon-gluon scattering, in addition to single- and double-diffractive interactions. The Monte Carlo-estimated cross section for the "hard" portion of these events, which includes the two-body interactions and double-diffractive events, is 47 mb. The lower-occupancy single-diffractive events have a cross section of approximately 12 mb, giving a total cross-section for the Pythia events of 59 mb. Based on very preliminary occupancy studies conducted in 2000 with pre-production versions of

the CFT electronics, we selected for these studies to overlay minimum bias events with a Poisson distribution averaging 7.5 Pythia events per beam crossing. Based on the above cross sections, we expect an average of 5 "hard (non-diffractive)" minimum bias events at an instantaneous luminosity of 5×10^{32} $cm^{-2}s^{-1}$. If 7.5 events representing the total cross section are used, this corresponds instead to approximately 6 "hard" minimum bias events, resulting in additional CFT occupancy. This increase in occupancy was necessary to match the occupancy measured with the pre-production electronics, but should be verified in the current system.

Recently, we have been able begin a detailed study of the CFT occupancy and signal-to-noise performance using the production electronics. Preliminary results are given in Figure 6, which shows the layer-by-layer occupancy of fibers in the CFT in single minimum bias events from Tevatron collisions, compared with that from the simulated Pythia minimum bias events. The occupancy in data was derived from minimum- and zero-bias trigger samples. The appropriate luminosity-weighting was used in subtracting the zero-bias occupancy to obtain the equivalent occupancy of a single minbias event. In order to make an appropriate comparison, only those Pythia events with pass a simulation of the minimum bias trigger used in the real data are considered. Finally, the Pythia occupancy has been scaled by a factor of 6/5 to represent the effective average occupancy of a single "hard" minimum bias event from our Monte Carlo sample. With these corrections, the agreement is quite good. We have indications that the discrepancy at low radius can be explained by detailed treatment of the material in the Monte Carlo and are working to resolve this.

3.3.3 Performance of the Run IIa Track Trigger

We have tested the functionality of the offline track trigger simulation by comparing its performance to that of the offline tracking. Using hits from Collider events where two muons have been reconstructed in the decay $Z\rightarrow\mu\mu$, we find a tracking efficiency of approximately 90% from the L1CTT for the high p_T muons. An example of this success is shown in Figure 7, where the left event display shows the tracks found by the trigger simulation overlaid on the tracks reconstructed offline. The event display on the right shows the locations of the two muons in the muon system.

In order to use the hits from the CFT, the correct track-finding equations were generated using the "as built" geometry of the CFT instead of a perfect detector geometry, thereby testing one of the crucial components of the Run IIa L1CTT and demonstrating the validity of the equation algorithms. Work is underway to understand the efficiency loss within the general context of commissioning the system.

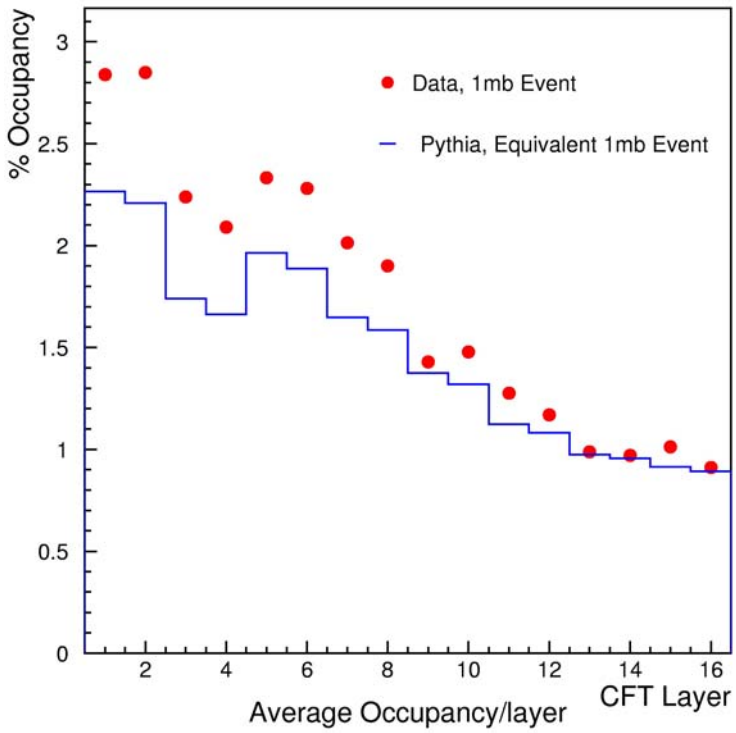

Figure 6. A comparison of the CFT occupancy by layer (the order is XUXV…) for minimum bias events in collider data and those generated by the Pythia Monte Carlo. See text for a detailed description.

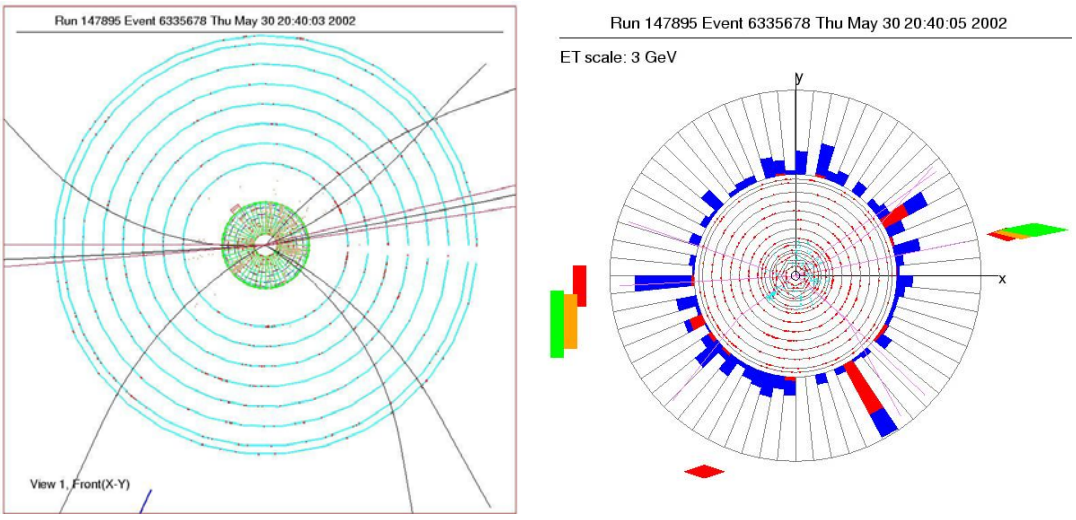

Figure 7. An example of the performance of the L1CTT using offline hits. The event display on the left shows the offline reconstructed tracks and the trigger sectors (The wedges at ~3 o'clock and ~9 o'clock) where the L1CTT simulation found high-p_T tracks. The display on the left shows the position of the muons found in the muon system at the same azimuth (the red/orange/green boxes).

3.4 Track Trigger Upgrade Strategy

As demonstrated above, the primary concern with the track trigger is the increase in rate for fake tracks as the tracker occupancy grows. Since the current track trigger requires hits in all 8 axial doublet layers, the only path to improving trigger rejection is to improve the trigger selectivity by incorporating additional information into the trigger algorithm. The short timescale until the beginning of Run IIb and resource limitations conspire to make it impossible to improve the physical granularity of the fiber tracker or to add additional tracking layers to the CFT. Instead, we propose to upgrade the track trigger logic. Essentially, more processing power allows the use of the full granularity and resolution of the individual CFT fibers rather than doublets in Level 1 track finding. Studies of this approach are presented below.

3.5 Singlet Equations in Track Finding

3.5.1 Concept

The motivation behind the use of singlet equations is illustrated in Figure 3, which shows a fragment of a CFT doublet layer. The thick black lines mark the area corresponding to a doublet hit, the current granularity of the L1CTT. As one can see from Figure 3, the doublet is larger than the fiber diameter. Since the hits from adjacent fibers are combined into the doublets before the tracking algorithm is run, this results in a widening of the effective width of a fiber to that of a doublet, decreasing the resolution of the hits that are used for track finding. In particular, the doublet algorithm is such that if fibers 1 and 4 shown on Figure 3 are hit, the trigger considers the doublet formed by fibers 1 and 2 *and* the doublet formed by fibers 3 and 4 to be hit. As the single-fiber occupancy grows, the application of this doublet algorithm results in a disproportionate increase in the hit occupancy seen by the trigger.

Hit patterns based instead on single fibers will be inherently narrower and will therefore have a reduced probability of selecting a random combination of hits. We have simulated different trigger configurations which include the all-singlet case (16 layers), as well as mixed schemes where some CFT layers are treated as pairs of singlet layers and the rest as doublets. In order to label the schemes we use the fact that the 8 layers of the CFT are labeled from A to H (see Figure 2). We use upper case letters to indicate that hits in this layer were treated as doublets; lower case letters indicate singlets. In this notation “ABCDEFGH” indicates the Run IIa CTT scheme with 8 layers of doublets and “abcdefgh” indicates 16 layers of singlets. Equations specifying which fibers should be hit as a function of momentum and azimuthal angle were generated for all configurations. Note that, in the results reported here, the equations have been generated specifying only which fibers should be hit, and not using vetoes on fibers that should not be hit. This will be discussed more completely in the next Section. , but it should be noted here that only stipulating the hit fibers insures that the efficiency of the track-finding algorithms is not compromised at high detector occupancies; if the hit fibers are present, the tracks will be found.

Because of the space-filling structure of the CFT shown in Figure 3, the number of fibers hit by a track passing through all 8 layers of the CFT varies with azimuthal position. This is shown in Figure 8, where the probability that a track will have ≥8, ≥10, ≥11, ≥12 and 13 hits out of 16 possible for the 16 layer singlet trigger scheme (abcdefgh) is plotted as a function of track sagitta. Here, it is assumed that fibers are 100% efficient.

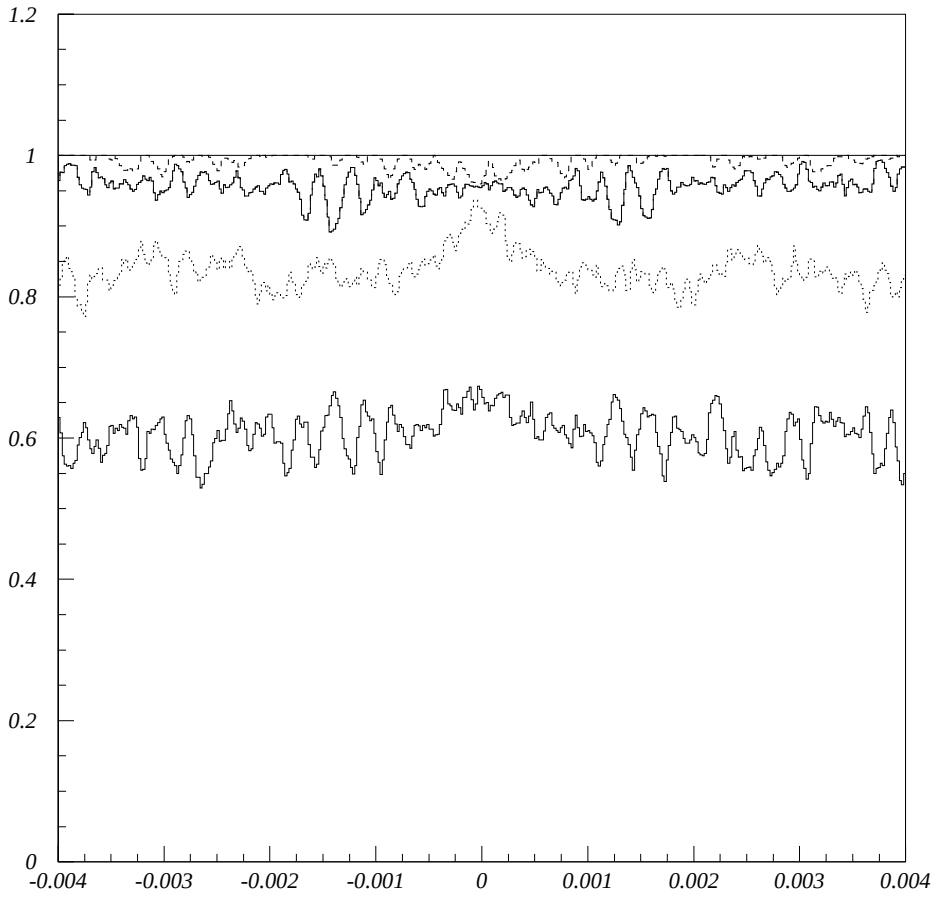

Figure 8. Geometrical acceptance for a charged particle to satisfy a ≥8 (solid line), ≥10 (dashed curve), ≥11 (upper solid curve), ≥12 (dot-dashed curve) and 13 (lower solid curve) hit requirement in the 16-trigger layer configuration "abcdefgh", versus the particle track sagitta, s = 0.02*e/ p_T, for a track starting at the center of a CFT trigger sector.

The baseline design uses a combination of all-singlet (abcdefgh) equations and combinations of singlets and doublets (abcdEFGH) depending on the momentum bin in question. The choice of algorithm in each case has been motivated by a consideration of the available FPGA resources, the rate of fake tracks at high occupancy, and the underlying physics goals for tracks at each momentum level. Tracks in the highest p_T bin will be found using the 16-layer

singlet equations (abcdefgh), which give the highest possible resolution and the best rejection of fake tracks, all of which is aimed at running a high-p_T isolated track trigger. In contrast, the lowest momentum bin (1.5 GeV<p_T<3 GeV) is covered by a hybrid scheme (abcdEFGH) with high granularity in the inner layers of the CFT and doublets in the outer layers where the occupancy is lower and multiple scattering will move the tracks off of their ideal trajectories. This choice is motivated by reducing the sheer number of equations, which increases as $1/p_T$, the desire to maintain high-efficiency for these tracks, since they are critical for *b*-tagging in the STT, and the relative insensitivity of the STT to some number of fake tracks. A summary of the chosen algorithms for each momentum bin is shown in Table 2.

3.5.2 Equation Algorithm Generation

An "equation" as defined here is a set of 12 or 16 terms specifying which fibers are hit by the tracks that traverse a given trajectory in a CFT Trigger sector. The equations are generated by a uniform sampling of all possible trajectories across a given trigger sector using many millions of iterations. The hit fibers along each trajectory are recorded. Trajectories with identical sets of hit fibers are combined, such that each equation corresponds to a specific average p_T and a specific location in ϕ. The relative "acceptance" of each equation is determined by calculating how many of the possible trajectories are represented by each of the equation terms.

The effects of CFT fiber inefficiencies/readout thresholds have been explicitly included in the equation generation. For each trajectory, the expected number of photoelectrons for each hit fiber is calculated from a Landau distribution based on the path length of the trajectory through the fiber. A threshold cut corresponding to 1.5 photoelectrons is then applied to each fiber in order to determine whether or not that fiber should appear as "hit" in the equation pattern. In this way, after many samplings, high efficiency is maintained on average for those trajectories which pass through the edges of fibers. Note that, since misses due to inefficiency and geometry are explicitly allowed, not all of the 12 or 16 terms need contain a fiber index.

Given the limits of finite FPGA resources, an optimization must be done separately for each of the p_T bins considered by the track trigger. Just as an example, if all of the singlet equations for the highest p_T bin (p_T>10 GeV), including those allowing some number of missed layers, were to be implemented in hardware, it would take some 300 times the FPGA resources of the current Run IIa system. Since this is clearly untenable, algorithms have been created to "prune" the sets of equations down to an acceptable level of complexity while maintaining high efficiency and low fake rates.

A simple set of cuts has been applied to reduce the number of equations needed in track finding:

1. All equations with no hit in any CFT doublet layer are discarded.

2. Equations with small relative acceptance of the total track phase space are discarded. For the 16-singlet equations, any single equation with a relative acceptance less than 3×10^{-5} of the total ensemble of possible trajectories is rejected. For the hybrid 12-layer scheme, only equations with an acceptance of 1.5×10^{-6} are kept.
3. Redundant equations having more hit fibers than a "shorter" overlapping equation are removed. A minimum of 8 hit fibers is required.

Motivations for these cuts are as follows. Hits in the middle layers of the CFT give the most precise information on the sagitta of the found tracks. Without a constraint in this region, real lower-momentum tracks are more easily "promoted" to fake high-p_T tracks with the addition of noise. As far as equation acceptance is concerned, many of the generated equations are satisfied by an extremely small number of possible tracks, but are as likely as any other equation to be satisfied by random noise hits. Therefore, the acceptance for fake tracks can be dramatically reduced by removing those equations that are essentially only efficient for picking up fake tracks due to noise. Removal of redundant equations is a simple technique: an equation specifying a trajectory that hit 8 specific fibers will still be satisfied by a more restrictive equation that hit 2 other specific fibers in addition to the original 8. As long as more fakes are not generated by the use of less-restrictive equations, full efficiency can be maintained with fewer equations. The high-p_T bin serves well as an example of the success of these methods: there are 697k original generated equations for the high-pT bin shown in Table 2. After the above cuts have been applied, only 9.4k equations remain, with little loss in efficiency and a much lower fake rate.

3.5.3 Rates and Rejection Improvements

The existing trigger simulation was adapted to make a realistic estimate of the trigger performance. Single muons were generated, overlaid on events containing a Poisson-distributed set of (Pythia) minimum bias interactions with a mean of 7.5 interactions per crossing and put through the detailed DØ simulation. (See the discussion of Monte Carlo generation, cross sections, and luminosity in Section 3.3.2.) They were then put through the modified trigger simulator. The fraction of events in which a trigger track matched the muon is defined as the trigger efficiency. A separate sample containing Poisson-distributed minimum bias events with the same 7.5 events/crossing rate is used to measure the fake rate, since there are no high-p_T tracks in this sample. All of the results presented in this section were obtained using a detailed description of the photons generated in the CFT. The ADC spectrum of the Monte Carlo hits was tuned to match that observed in the data, and a realistic threshold cut equivalent to 1.5 photoelectrons was applied to each of the channels.

The results of the procedure described are summarized in Table 2, which shows the performance of the baseline L1CTT algorithmic design. Shown in the table for each of the momentum bins is the singlet/doublet scheme employed, the fraction of each of the pure 7.5 minbias events that produces a fake track in each

momentum range and the number of equations used in each trigger sector. One point deserves mention here: the fake rate in the lowest p_T bin seems very high, except that huge numbers of low-p_T tracks are present at such high luminosities. The actual ratio of fake low-p_T tracks to real ones present in the Monte Carlo is approximately 1:1. This is well within the rate of tracks that can be handled comfortably by the STT and thus this algorithm will not compromise *b*-tagging at high-luminosity.

Table 2. Performance of the baseline L1CTT upgrade. The tracking efficiency, rate of fake tracks, and resources required by each of the momentum bin algorithms in the upgraded L1CTT. These quantities were evaluated using Pythia Monte Carlo events with 7.5 Poisson-distributed minimum bias interactions. Track-finding efficiency is based on single muon events with background interactions overlaid. For a comparison, the default Run IIa L1CTT "ABCDEFGH" is shown for the high-p_T bin only. The rate of fake tracks is the fraction of events in which one or more fake tracks is generated in the specified momentum bin. Resources are defined by the number of equations needed for each momentum bin (the first figure) times (X) the number of terms (= the number of layers) in each equation set. The superior performance of the new algorithms is evident.

Scheme/p_T Range	Track Finding Efficiency (%)	Rate of Fake Tracks (% of events)	Resources
ABCDEFGH (p_T >10)	96.9	1.03 $\pm$ 0.10	11k X 8
abcdefgh (p_T >10)	98.03 $\pm$ 0.22	0.056 $\pm$ 0.009	9.4k X 16
abcdEFGH ($5 < p_T < 10$)	99.20 $\pm$ 0.14	0.89 $\pm$ 0.11	8.9k X 12
abcdEFGH ($3 < p_T < 5$)	98.40 $\pm$ 0.20	4.5 $\pm$ 0.2	11.3k X 12
abcdEFGH ($1.5 < p_T < 3$)	95.15 $\pm$ 0.32	25.4 $\pm$ 0.2	15.5k X 12

As can clearly be seen from the table, the performance of the trigger using the singlet resolutions is far superior to the default Run IIa L1CTT installation. In the high p_T region, the fake rate has decreased by a factor of 20 and the efficiency is slightly higher. Given that the L1CTT serves as a basis for muon and, in the upgrade design, calorimeter triggers, such an improvement in performance gives us confidence that the L1CTT will be a viable trigger at the highest luminosities of Run IIb. The discussion of the implications of the FPGA resources will be presented within the hardware discussion, below.

3.6 Implementation of Level 1 Central Track Trigger

3.6.1 CTT Electronics

Figure 9 shows the block diagram of the existing L1 central track trigger electronics. The track trigger hardware has three main functional elements.

The first element consists of the Analog Front-End (AFE) boards that receive signals from the VLPCs. The AFE boards provide both digitized information for L3 and offline analysis as well as discriminated signals used by the CTT. Discriminator outputs for 128 channels are buffered and transmitted over a fast link to the next stage of the trigger. The axial layers of the CFT are instrumented using 76 AFE boards, each providing 512 channels of readout. The axial CPS strips are instrumented using 10 AFE boards, each having 256 channels devoted to axial CPS readout and the remaining 256 channels devoted to stereo CFT readout. The FPS is instrumented using 32 AFE boards. Additional AFE boards provide readout for the stereo CPS strips and remaining stereo CFT fibers.

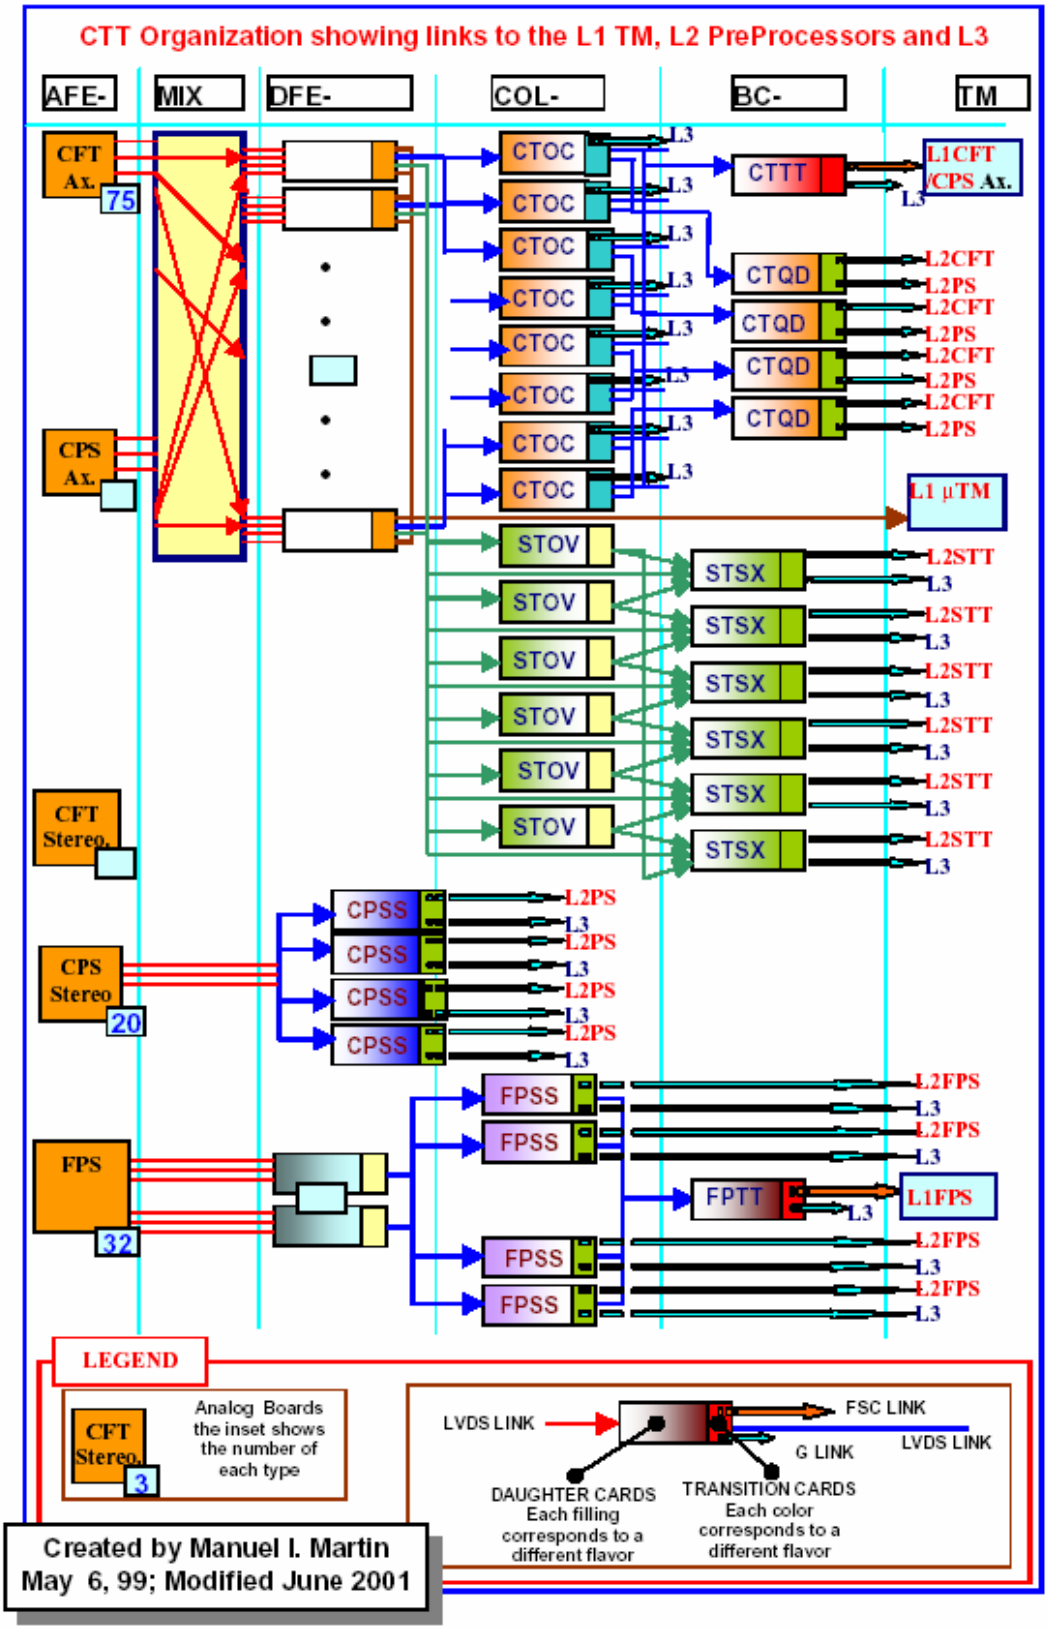

Figure 9. Block diagram of level 1 central track trigger.

The second hardware element is the Mixer System (MIX). The MIX resides in a single crate and is composed of 20 boards. It receives the signals from the AFE boards and sorts them for the following stage. The signals into the AFE boards are ordered in increasing azimuth for each of the tracker layers, while the trigger is organized into TS wedges covering all radial CFT/CPS axial layers within 4.5 degrees in ϕ. Each MIX board has sixteen CFT inputs and one CPS input. It shares these inputs with boards on either side within the crate and sorts them for output. Each board then outputs signals to two DFEA boards (described below), with each DFEA covering two TS.

The third hardware element is based on the Digital Front-End (DFE) motherboard. These motherboards provide the common buffering and communication links needed for all DFE variants and support two different types of daughter boards, single-wide and double-wide. The daughter boards implement the trigger logic using FPGA chips. These have a very high density of logic gates and lend themselves well to the track equations. Within these chips all 17k equations are processed simultaneously in under 200 ns. This design also keeps the board hardware as general as possible. The motherboard is simply an I/O device and the daughter boards are general purpose processors. Since algorithms and other details of the design are implemented in the FPGA, which can be reprogrammed via high level languages, one can download different trigger configurations for each run or for special runs and the trigger can evolve during the run. The signals from the Mixer System are received by 40 DFE Axial (DFEA) boards. There are also 5 DFE Stereo (DFES) boards that prepare the signals from the CPS stereo layers for L2 and 16 DFEF boards that handle the FPS signals.

3.6.2 CTT Outputs

The current tracking trigger was designed to do several things. For the L1 Muon trigger it provides a list of found tracks for each crossing. For the L1 Track Trigger it counts the number of tracks found in each of four p_T bins. It determines the number of tracks that are isolated (no other tracks in the home sector or its neighbors). The sector numbers for isolated tracks are recorded to permit triggers on acoplanar high p_T tracks. Association of track and CPS clusters provides the ability to recognize both electron and photon candidates. FPS clusters are categorized as electrons or photons, depending on an association of MIP and shower layer clusters. Finally, the L1 trigger boards store lists of tracks for each beam crossing, and the appropriate lists are transferred to L2 processors when an L1 trigger accept is received.

The L1 CTT must identify real tracks within four p_T bins with high efficiency. The nominal p_T thresholds of the bins are 1.5, 3, 5, and 10 GeV. The L1 CTT must also provide rejection of fake tracks (due to accidental combinations in the high multiplicity environment). The trigger must perform its function for each beam crossing at either 396 ns or 132 ns spacing between crossings. For each crossing a list of up to six found tracks per p_T bin is packed into 96 bits and transmitted from each of the 80 trigger sectors. These tracks are used by the L1

muon trigger and must be received within 1000 ns of the crossing. These track lists are transmitted over copper serial links from the DFEA boards.

The L1 CTT counts the number of tracks found in each of the four p_T bins, with subcategories such as the number of tracks correlated with showers in the Central Preshower Detector, and the number of isolated tracks. Azimuthal information is also preserved so that information from each ϕ region can be correlated with information from other detectors. The information from each of the 80 TS is output to a set of 8 Central Tracker Octant Card (CTOC) boards, which are DFE mother boards equipped with CTOC type double-wide daughter boards. During L1 running mode, these boards collect the information from 10 DFEA boards, combine the information, and pass it on to a single Central Track Trigger Terms (CTTT) board. The CTTT board, also a DFE-type mother board equipped with a similar double wide daughter board, assembles the information from the eight CTOC boards and makes all possible trigger terms for transmission to the Trigger Manager (TM). The TM constructs the 32 AND/OR terms that are used by the Trigger Framework in forming the L1 trigger decision. For example, the term “TPQ(2,3)” indicates two tracks associated with CPS hits were present in quadrant 3. Additional AND/OR terms provide CPS and FPS cluster characterization for use in L1. The Trigger Framework accommodates a total of 256 such terms, feeding them into a large programmable AND/OR network that determines whether the requirements for generating a trigger are met.

The DFEA boards store lists of tracks from each crossing, and these lists are transferred to the L2 processors when an L1 trigger accept is received. A list of up to 6 tracks is stored for each p_T bin. When an L1 trigger accept is received, the normal L1 traffic is halted and the list of tracks is forwarded to the CTOC board. This board recognizes the change to L2 processing mode and combines the many input track lists into a single list that is forwarded to the L2 processors. Similar lists of preshower clusters are built by the DFES and DFEF boards for the CPS stereo and FPS strips and transferred to the L2 processors upon receiving an L1 trigger accept.

3.6.3 Implementation of the Run IIb upgrade

The implementation, cost and schedule depends largely on the algorithm chosen and what FPGA resources the algorithm requires. The entire track finding logic is included on the 80 DFEA daughter boards located on 40 mother boards. The proposed upgrade will be restricted to replacing the existing DFEA daughterboards with new ones. The current design already brings the singlet hits onto these daughter boards so we do not anticipate any changes to the mother boards or any of the upstream electronics and cabling. The current system publishes the six highest p_T tracks in each of four momentum bins (24 tracks). The new design will do the same so that no changes are needed to the output daughter cards nor to the downstream cabling. The crate controller card stores the logic in flash memory for local downloading to the other cards in the crate. The present design uses 1.6 Mbytes and the controller can hold 512 Mbytes,

giving an expansion factor of more than 250. Larger gate array chips do not use much more power so that the power supplies and cooling are also adequate.

3.6.4 DFE motherboard

The Digital Front End (DFE) Motherboard is a general purpose, high bandwidth platform for supporting reconfigurable logic such as FPGAs. It is intended for applications where a DSP or other microprocessor is too slow. The DFE motherboard is a 6U x 320mm Eurocard with fully custom hard metric backplane connectors.

Ten point-to-point links bring data onto the DFE motherboard at an aggregate data rate of 14.8 Gbps. The physical link consists of five twisted pairs (Low Voltage Differential Signals) and is terminated with hard metric female connectors on the front panel of the DFE motherboard. After entering the DFE motherboard, the ten links are sent to receivers, which convert the serial data back to a 28 bit wide bus running at 53 MHz. These busses, in turn, are buffered and routed to the two daughtercards. Links 0, 1, and 2 are sent to the top daughtercard; links 7, 8, and 9 are sent to the bottom daughtercard; and links 3, 4, 5, and 6 are sent to both the top and bottom daughtercards. A basic dataflow diagram of the motherboard is shown in Figure 10.

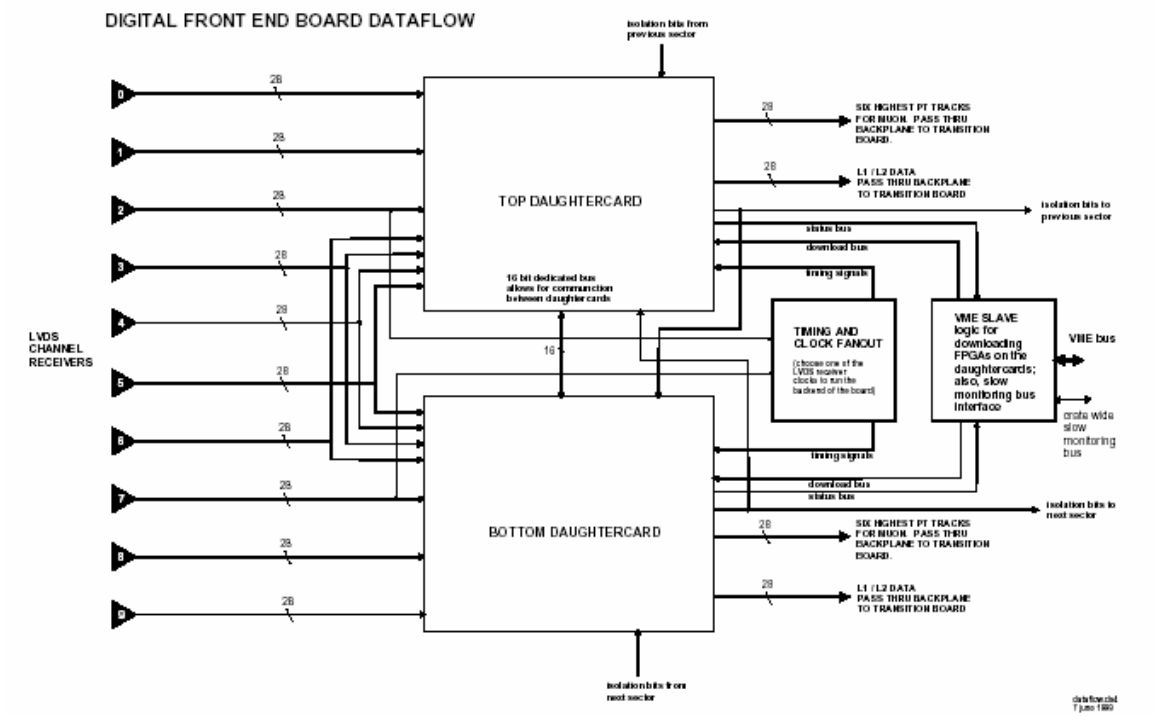

Figure 10. Block diagram of DFE motherboard.

The outputs from the daughterboards are fed back to the motherboard and are passed to hard metric connectors through the backplane to a transition card. The transition card converts the output busses back into LVDS channel links to feed the output from one DFE motherboard into a second DFE motherboard that is reconfigured as a data concentrator.

One of the main purposes of the DFE motherboard is to support programmable logic contained on the daughterboards. Daughterboard PLDs are exclusively FPGAs, which programmable logic on the daughtercards is an FPGA, which needs to be downloaded after each power cycle, since its configuration memory is volatile. As the density of FPGAs increases, so does the size of the configuration data file that must be downloaded to it. The size of the configuration data files may be several megabytes per daughtercard. For this reason, a custom high speed bus is incorporated into the DFE backplane. Slot 1 of the backplane is reserved for a custom DFE crate controller.

3.6.5 DFEA daughterboard

The DFEA daughterboard is a 10-layer PC board, 7.8" x 4.125" in size. It has 500 gold-plated contacts on the "solder" side to mate with spring loaded pins located on the DFE motherboard. Sixteen bolts attach each daughterboard to the motherboard. Figure 11 shows a photograph of the Run IIa DFEA daughterboard. The motherboard supports two DFEA daughterboards.

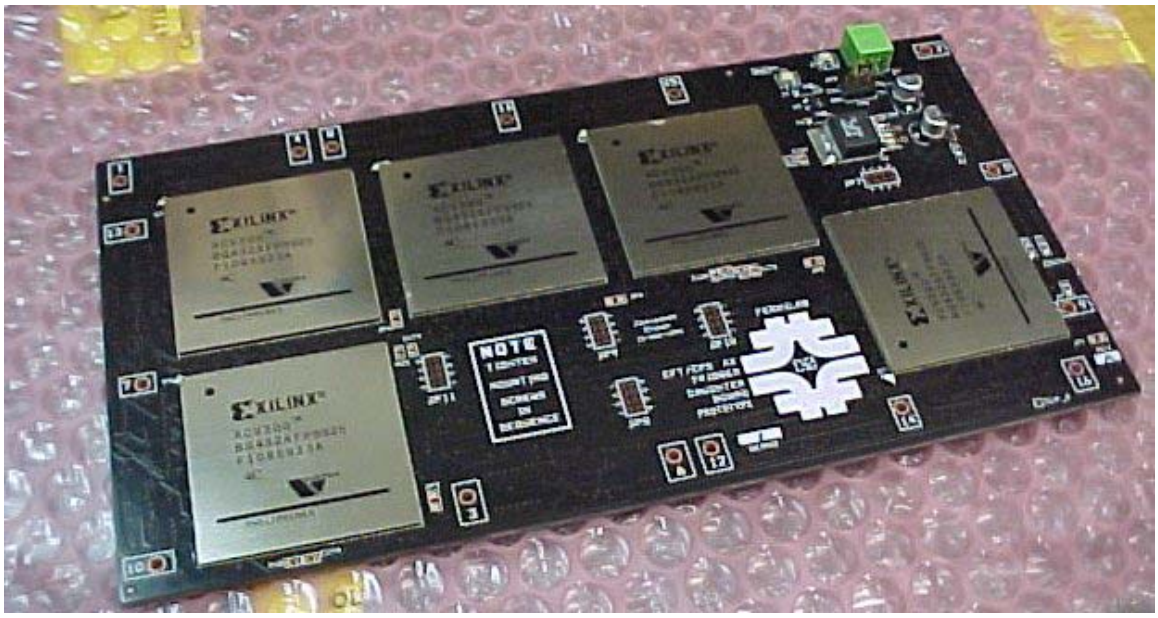

Figure 11. Photograph of Run IIa DFEA daughterboard.

Each DFEA daughterboard must perform the following functions (see also the block diagram in Figure 12):

- Find tracks in each of four p_T bins (Max, High, Med, and Low).
- Find axial CPS clusters.
- Match CPS clusters and tracks.
- Count tracks and clusters (matched, isolated, and non isolated) for L1 readout.
- Store tracks and clusters for L2 readout.
- Generate a list of the six highest p_T tracks to send to Muon L1

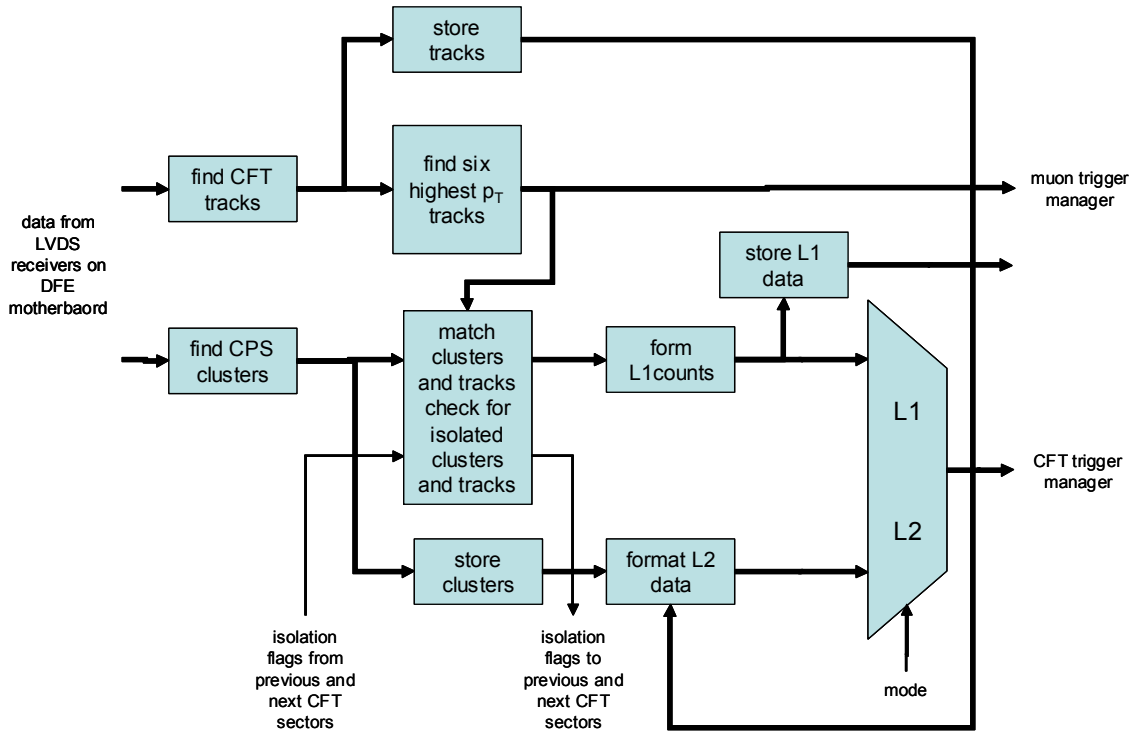

Figure 12. Block diagram of DFEA daughterboard functionality.

Track finding is the most difficult and expensive function of the DFEA daughterboard. To identify tracks down to 1.5 GeV, relatively large FPGAs must be used. These FPGAs match the raw data to a predefined list of track equations and serialize the found tracks to be read out at 53 MHz. The present daughter board houses five Xilinx Virtex-I chips. These consist of one Virtex 600, three Virtex 400's and one Virtex 300. They are housed in pin ball grid array packages. The PC board requires 10 layers to interconnect these chips. The present Virtex 600 has an array of 64x96 slices with each slice containing 2 four-input look up tables (LUT) giving a total of 12,288 LUT's. Figure 13 shows a block diagram of these FPGAs.

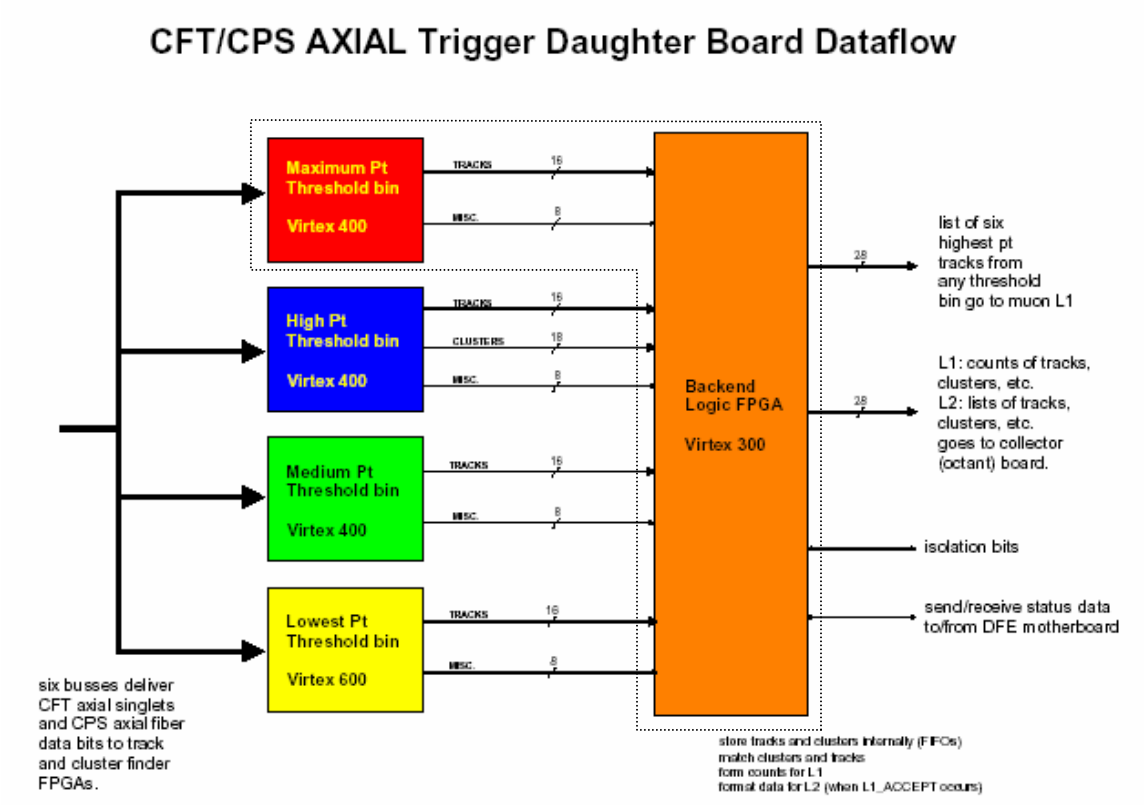

Figure 13. Block Diagram of the five FPGAs on the Run IIa DFEA daughterboard and their interconnection. For the Run IIb upgrade four larger FPGAs will replace the four FPGAs on the left and the function of the backend logic FPGA will be incorporated into one of these four FPGAs as shown by the dotted outline.

Inside each track finder FPGA the fiber bits are matched to the pre-defined track equations in parallel (combinatorial logic). The output is a 44x8 or 44x16 matrix of bits. A sequential pipeline performs a priority encode over this matrix and reports the six highest p_T tracks in each chip (http://d0server1.fnal.gov/users/jamieson/www/notes/2001-12-29a.pdf). This is shown schematically in Figure 14. In addition to large combinatorial logic resources, each of these FPGAs has to accommodate a large number of inputs. This results in a device that requires use of BGA (Ball Grid Array) packages.

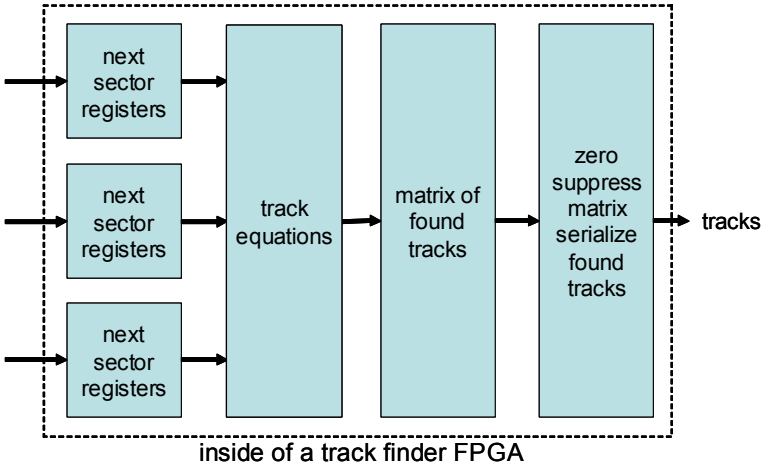

Figure 14. Block diagram of one track finder FPGA.

The algorithms that we are considering for Run IIb require at least 10 times more resources than Run IIa algorithm, which the present daughter boards cannot provide. Xilinx does not intend to produce more powerful chips that are pin-compatible with the Virtex-I chips we are using now. If we want to use newer, more powerful FPGAs, new daughter boards have to be designed.

The Virtex-II series FPGAs have 8 to 10 times larger logic cells than the largest chips that we are currently using. The Virtex-II series offers chips which have 2M to 8M system gates and about 25K to 100K logic cells (see Table 3). These chips come in a ball grid array packages similar in size to the existing parts. Thus, we will be able to fit four of these chips on new daughter boards of the same size as the present daughter boards. Due to the denser parts the PC boards may require 2 or 4 additional layers.

In addition, the speed of the Virtex-II chips is in the range of 200-300 MHz. We are looking into the gains we may achieve by utilizing this increased speed and similarities of “sub-units” of different equations. By running the chips at higher speeds, we may be able pipeline some of the processing allowing possible reuse of similar “sub-units” of equations stored in different logic cells, and therefore accommodate a larger number of equations.

Table 3. Virtex-II Field Programmable Gate Array Family Members.

Device	System Gates	CLB (1 CLB = 4 slices = Max 128 bits)			Multiplier Blocks	SelectRAM Blocks		DCMs	Max I/O Pads(1)
		Array Row x Col.	Slices	Maximum Distributed RAM Kbits		18-Kbit Blocks	Max RAM (Kbits)		
XC2V40	40K	8 x 8	256	8	4	4	72	4	88
XC2V80	80K	16 x 8	512	16	8	8	144	4	120
XC2V250	250K	24 x 16	1,536	48	24	24	432	8	200
XC2V500	500K	32 x 24	3,072	96	32	32	576	8	264
XC2V1000	1M	40 x 32	5,120	160	40	40	720	8	432
XC2V1500	1.5M	48 x 40	7,680	240	48	48	864	8	528
XC2V2000	2M	56 x 48	10,752	336	56	56	1,008	8	624
XC2V3000	3M	64 x 56	14,336	448	96	96	1,728	12	720
XC2V4000	4M	80 x 72	23,040	720	120	120	2,160	12	912
XC2V6000	6M	96 x 88	33,792	1,056	144	144	2,592	12	1,104
XC2V8000	8M	112 x 104	46,592	1,456	168	168	3,024	12	1,108

3.6.6 Resource evaluation

We estimated the resources required to implement 9.5k track equations in the highest p_T bin (p_T>10GeV) by performing a full simulation using the existing Run IIa firmware infrastructure implemented in VHDL. We modified the firmware to adapt to the different scheme proposed for Run IIb. Since we will use all 16 singlet layers in this p_T bin, it was necessary to eliminate the doublet-former module from the Run IIa firmware. The “flattened” fiber hit inputs are then sent directly to the track equation evaluator module. This module compares the fiber hits to the above set of equations to find a valid trigger track. This implementation also preserves the output structure of the firmware and the triggered track results are reported in terms of the matrix used for serialization of the tracks downstream.

For the highest p_T bin (p_T > 10 GeV), the device utilization report for the proposed implementation using the ISE synthesis tool available from Xilinx for firmware implementation and simulation is given in Table 4. We find that implementing 9.4K equations will utilize 11863 slices of the FPGA. This translates to 35% of a Virtex-II series XC2V6000 FPGA.

Table 4. ISE device utilization summary for XC2V6000.

Number of External IOBs*	122 out of 684	17%
Number of LOCed External IOBs*	0 out of 122	0%
Number of SLICEs	11863 out of 33792	35%

*IOB = input/output block

We have also implemented the 7.5K equations required by the lowest pT bin and find that we need 32% of the slices in XC2V6000 FPGA. In addition we have verified that the number of FPGA slices needed is proportional to the number of track equations. We are therefore confident that the resources required for intermediate bins scale with the number of track equations. Our studies show that the number of track equations for the four p_T bins will range from about 7.5K to 10K equations. Therefore we estimate that one XC2V6000 chip will be needed for each of the four p_T bins. The functionality of the backend FPGA which reports the final track trigger results to the downstream boards can be absorbed in one of the FPGA for the medium p_T bins. Figure 15 displays a drawing showing the footprints of four XC2V6000 chips on a DFEA daughterboard.

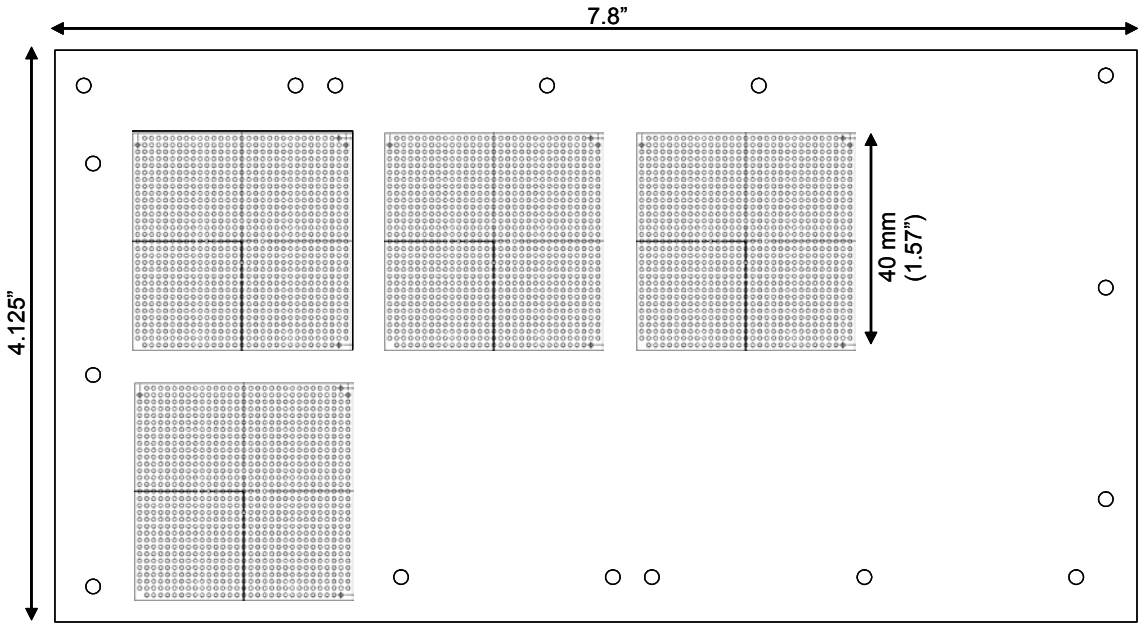

Figure 15. Drawing showing the footprints of four XC2V6000 chips on a DFEA daughterboard.

3.7 L1 Tracking Trigger Summary and Conclusions

Based upon current simulation results, it is clear that the L1 CTT needs to be upgraded in order to maintain the desired triggering capabilities as a result of the anticipated Run IIb luminosity increases. Because of the tight timescales and limited resources available to address this particular challenge, significant alterations to the tracking detector installed in the solenoid bore are not considered feasible.

Improving the resolution of the L1 CTT by treating CFT axial layers as singlets rather than doublet layers in the L1 trigger significantly improves the background rejection. Simulation studies show a factor of ten improvement in fake rejection rate at high-p_T by treating the hits from fibers on all axial layers as singlets.

The proposed FPGA upgrade provides a major increase in the number of equations and number of terms per equation that can be handled, and provides increased flexibility in the track finding algorithms that may be implemented. Depending on the p_T range, either mixtures of doublet and singlet layers or full singlet layers are proposed. Finally, we have demonstrated the technical feasibility of the upgrade by implementing the proposed algorithm in currently available FPGAs (e.g. Xilinx Virtex II series).

4 Level 1 Calorimeter Trigger

4.1 Goals

The primary focus of Run IIb will be the search for the mechanism of electroweak symmetry breaking, including the search for the Higgs boson, supersymmetry, or other manifestations of new physics at a large mass scale. This program demands the selection of events with particularly large transverse momentum objects. The increase in luminosity (and thus increasing multiple interactions), and the decreased bunch spacing (132ns) for Run IIb will impose heavy loads on the Level 1 (L1) calorimeter trigger. The L1 calorimeter trigger upgrade should provide performance improvements over the Run IIa trigger system to allow increased rejection of backgrounds from QCD jet production, and new tools for recognition of interesting signatures. We envision a variety of improvements, each of which will contribute to a substantial improvement in our ability to control rates at the L1 trigger. In the following sections we describe how the L1 calorimeter trigger upgrade will provide

- An improved capability to correctly assign the calorimeter energy deposits to the correct bunch crossing via digital filtering
- A significantly sharper turn-on for jet triggers, thus reducing the rates
- Improved trigger turn-on for electromagnetic objects
- The ability to make shape and isolation cuts on electromagnetic triggers, and thus reducing rates
- The ability to match tracks to energy deposition in calorimeter trigger towers, leading to reduced rates
- The ability to include the energy in the intercryostat region (ICR) when calculating jet energies and the missing ET
- The ability to add topological triggers which will aid in triggering on specific Higgs final states.

The complete implementation of all these improvements will provide us with the ability to trigger effectively with the calorimeter in the challenging environment of Run IIb.

4.2 Description of Run IIa Calorimeter Electronics

4.2.1 Overview

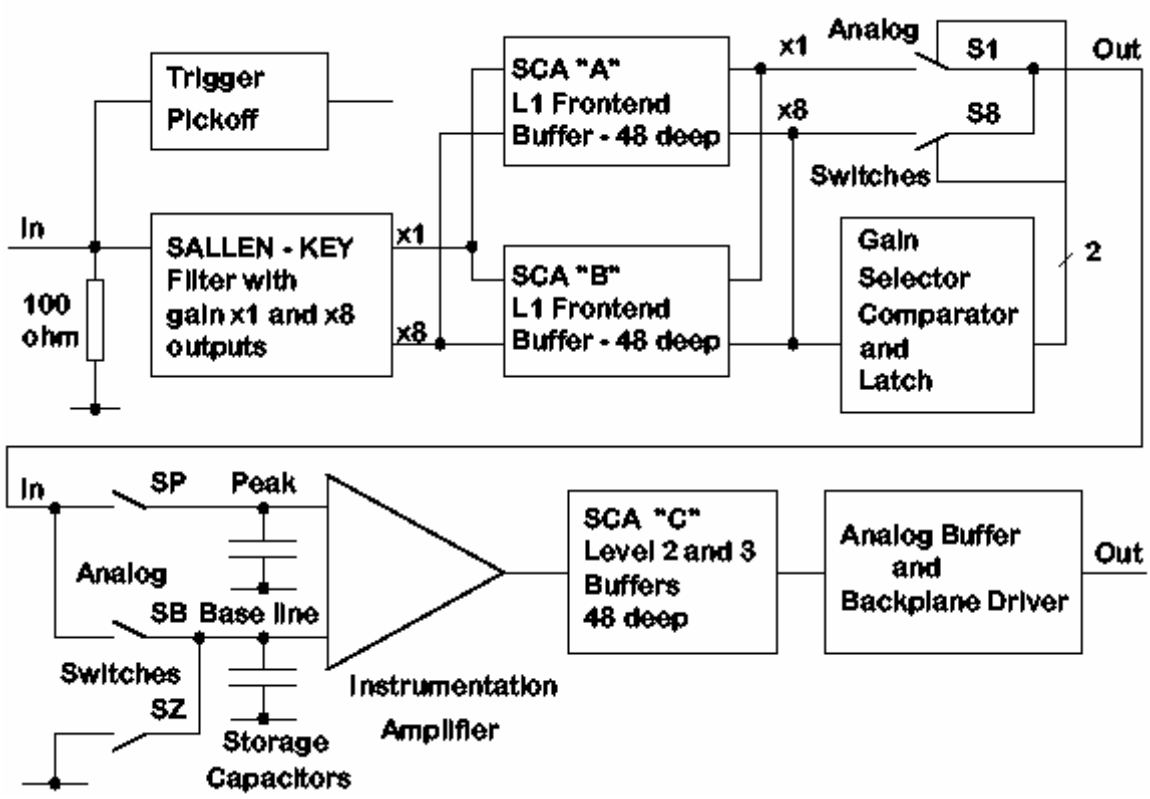

Figure 16. Functional diagram of the BLS system showing the precision readout path and the location of the calorimeter trigger pickoff signal.

The charge from the calorimeter is integrated in the charge sensitive preamplifiers located on the calorimeter. The preamplifier input impedance is matched to the 30 Ω coaxial cables from the detector (which have been equalized in length), and the preamplifiers have been compensated to match the varying detector capacitances, so as to provide signals that have approximately the same rise time (trace #1 in Figure 17). The fall time for the preamp signals is 15 μs. The signals are then transmitted (single ended) on terminated twisted-pair cable to the baseline subtractor cards (BLS) that shape the signal to an approximately unipolar pulse (see Figure 16 for a simple overview). The signal on the trigger path is further differentiated by the trigger pickoff to shorten the pulse width, leading to a risetime of approximately 120 ns (trace #2 in Figure 17). The signals from the different depths in the electromagnetic and hadronic sections are added with appropriate weights to form the analog trigger tower sums. These analog sums are output to the L1 calorimeter trigger after passing through the trigger sum drivers. The signals are then transported differentially (on pairs of 80Ω coaxial cable) ~80m to the L1 calorimeter trigger (the negative side of a differential pair is shown in trace #4 in Figure 17). The key elements of the calorimeter trigger path are described in more detail in the following sections.

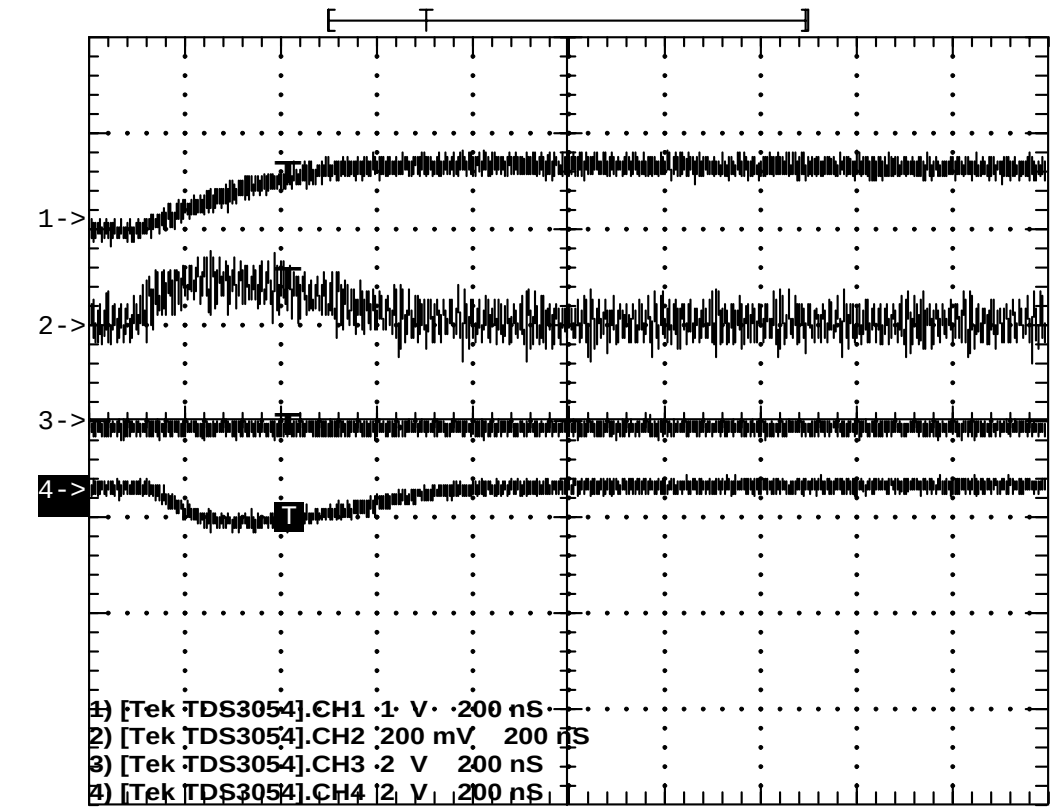

Figure 17. Scope traces for actual detector signals for an EM section. The horizontal scale is 200ns/large division. The top trace (#1, 1V/div) is of a preamp output signal as seen at the input to the BLS. The second trace (#2, 200mV/div) is of the trigger pickoff output on the BLS card (the large noise is due to scope noise pickup, so is not real). The fourth trace (#4, 2V/div) is the negative side of the differential trigger sum driver signal at the BLS that is sent to the L1 calorimeter trigger.

4.2.2 Trigger pickoff

The trigger pickoff captures the preamplifier signal before any shaping. A schematic of the shaping and trigger pickoff hybrid is shown in Figure 18 (the trigger pickoff section is in the upper left of the drawing). The preamplifier signal is differentiated and passed through an emitter follower to attempt to restore the original charge shape (a triangular pulse with a fast rise and a linear fall over 400 ns). This circuitry is located on a small hybrid that plugs into the BLS motherboard. There are 48 such hybrids on a motherboard, and a total of 55,296 for the complete detector.

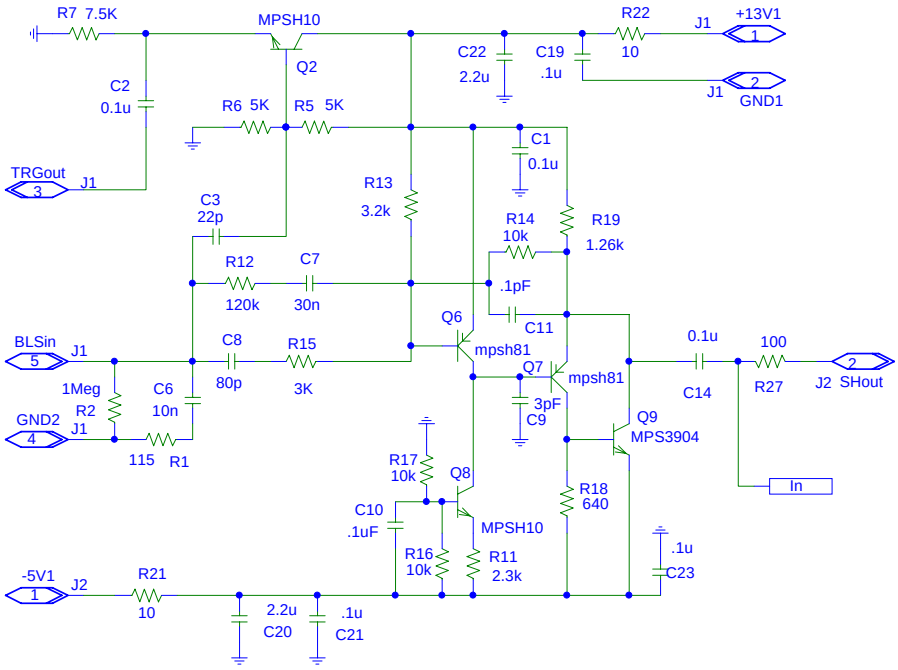

Figure 18. Schematic of the trigger shaper and trigger pickoff (upper left of picture). Pin 5 is the input, pin 3 is the trigger pickoff output, and pin 2 is the shaped precision signal output.

4.2.3 Trigger summers

The trigger pickoff signals for EM and HAD sections in individual towers (note these are not the larger trigger towers) are routed on the BLS board to another hybrid plug-in that forms the analog sums with the correct weighting factors for the different radial depth signals that form a single tower. The weighting is performed using appropriate input resistors to the summing junction of the discrete amplifier. A schematic for this small hybrid circuit is shown in Figure 19.

A single 48 channel BLS board has 8 trigger summer hybrids (4 EM towers and 4 HAD towers). There are a total of 9,216 hybrid trigger summers made up of 75 species. Since they are relatively easy to replace, changes to the weighting schemes can be considered. Recall, however, that access to the BLS cards themselves requires access to the detector as they are located in the area directly beneath the detector, which is inaccessible while beam is circulating.

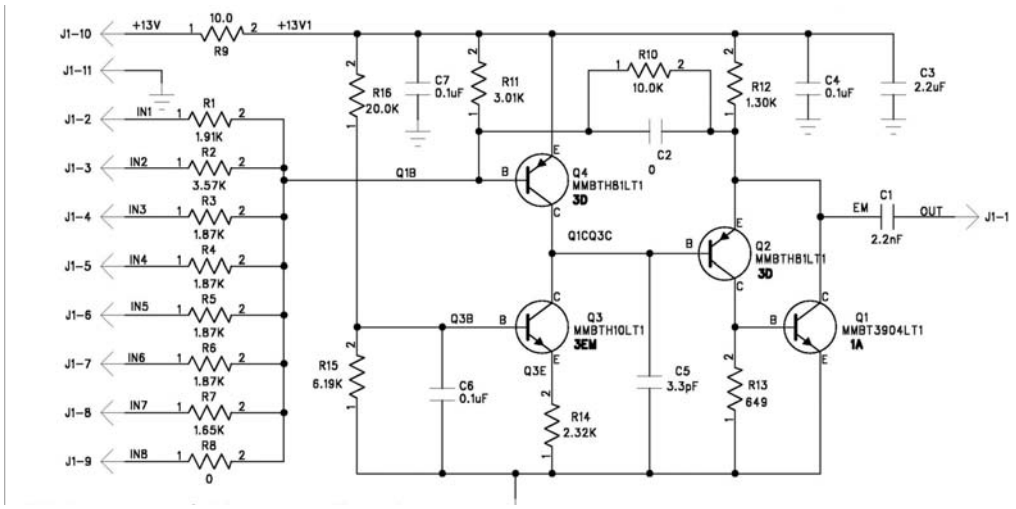

Figure 19. Schematic of the trigger summer hybrid. Up to 8 inputs from the various layers in a single tower can be summed with varying gains determined by the resistors to the summing junction (shown at left).

4.2.4 Trigger sum driver

The outputs of the 4 EM trigger summers and the 4 HAD trigger summers on a single BLS board are summed separately (except at high η) once more by the trigger sum driver circuit (see the schematic in Figure 20) where a final overall gain can be introduced. This circuit is also a hybrid plug-in to the BLS board and is thus easily replaceable if necessary (with the same access restrictions discussed for the trigger summers). In addition the driver is capable of driving the coaxial lines to the L1 Calorimeter trigger. There are a total of 2,560 such drivers in 8 species (although most are of two types).

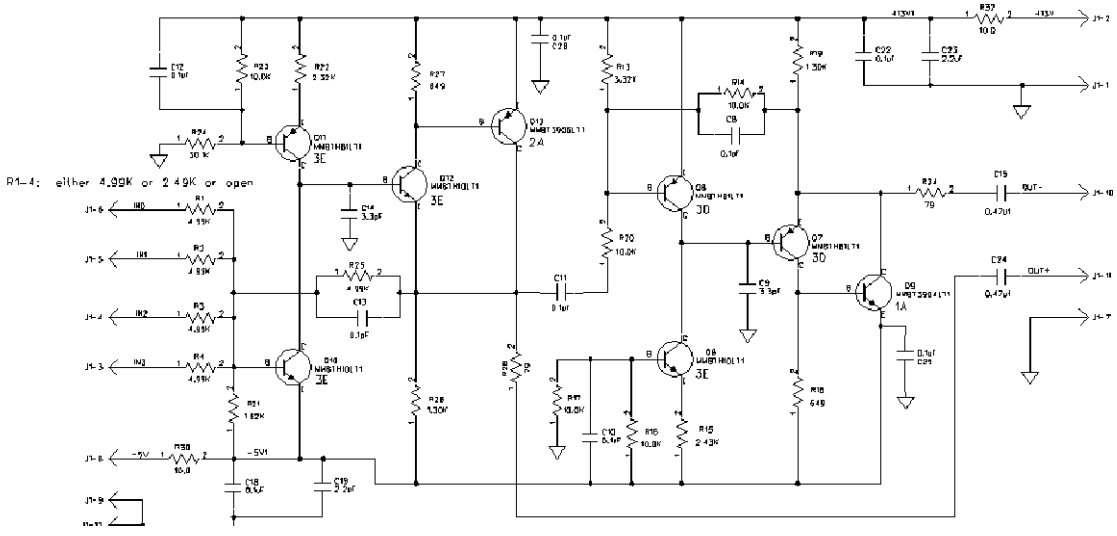

Figure 20. Schematic of the trigger sum driver hybrid. This circuit sums the outputs of up to 4 trigger summer outputs of the type shown in Figure 19.

4.2.5 Signal transmission, cable dispersion

The signals from the trigger driver circuits are transmitted differentially on two separate miniature coax (0.1") cables. The signal characteristics for these cables are significantly better than standard RG174 cable. However first indications are that the signals seen at the end of these cables at the input to the L1 calorimeter trigger are somewhat slower than expected (an oscilloscope trace of such a signal is shown in Figure 21 for EM and Figure 22 for HAD). The cause of the deviation from expectations is not presently known and is under investigation. It is possible that the signal dispersion in these coaxial cables is worse than expected. In any case, we must deal with these pulses that are over 400ns wide (FWHM) and thus span a few 132ns bunch crossings. The most effective treatment of this problem is to further process the signal through digital filtering to extract the proper bunch crossing. This solution is described in more detail in later sections.

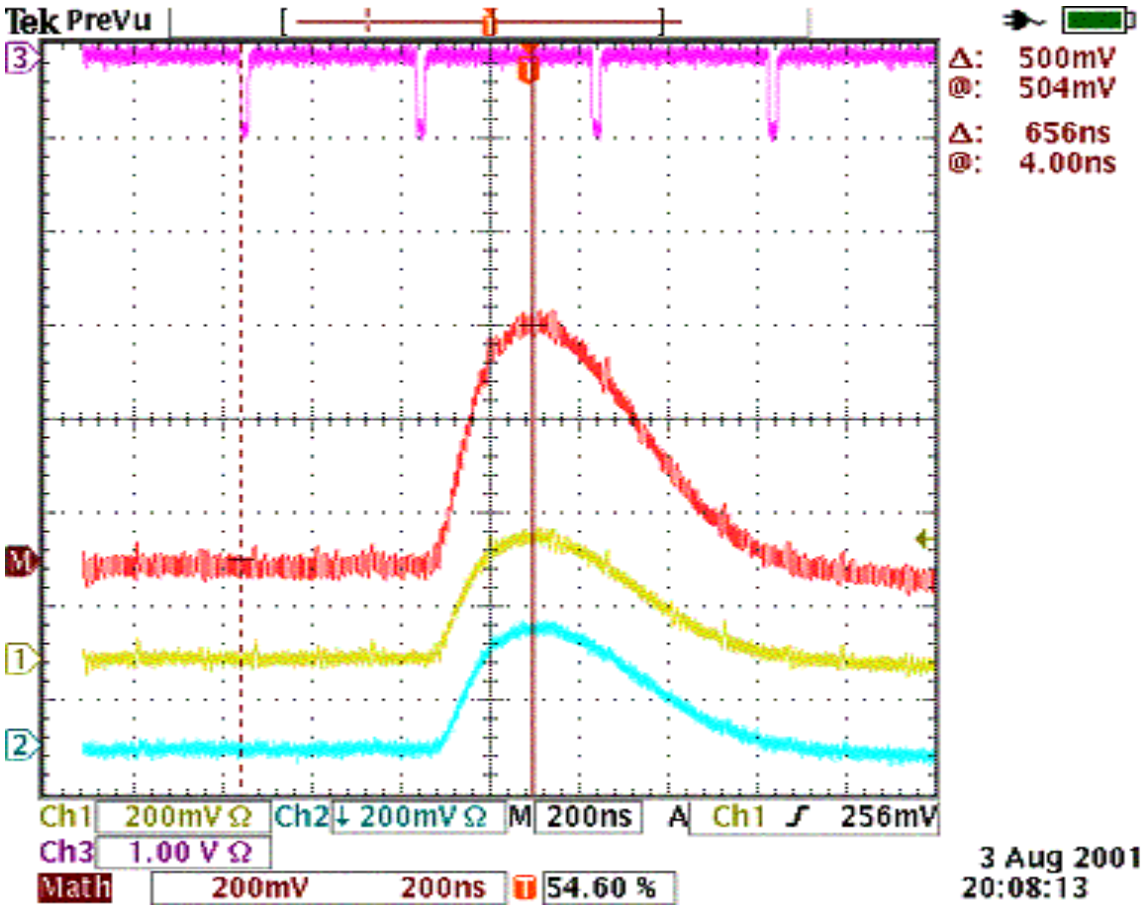

Figure 21. Actual traces of EM trigger tower (ieta=+1, iphi=17) data from the trigger sum driver signal as measured at the input to the L1 calorimeter trigger. The top trace (#3) shows the time of the beam crossings (396ns). The second trace (M) shows the addition of the two differential signals after inversion of the negative one. The third trace (#1) is the positive side of the differential pair. The fourth trace (#2) is the inverted trace for the negative side of the differential pair.

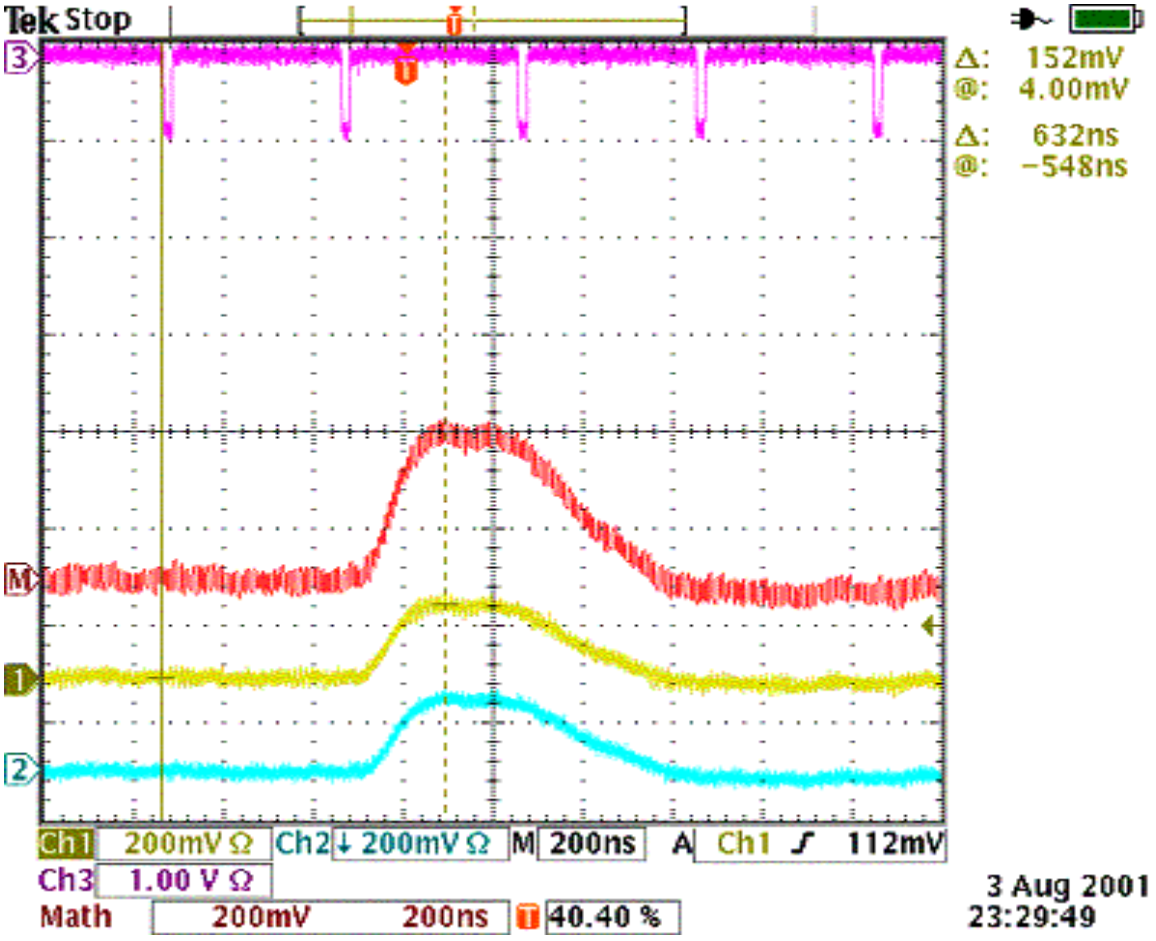

Figure 22. Actual traces of HAD trigger tower (ieta=+1, iphi=17) data from the trigger sum driver signal as measured at the input to the L1 calorimeter trigger. The top trace (#3) shows the time of the beam crossings (396ns). The second trace (M) shows the addition of the two differential signals after inversion of the negative one. The third trace (#1) is the positive side of the differential pair. The fourth trace (#2) is the inverted trace for the negative side of the differential pair.

4.3 Description of Current L1 Calorimeter Trigger

4.3.1 Overview

The DØ uranium-liquid argon calorimeter is constructed of projective towers covering the full 2π in the azimuthal angle, ϕ, and approximately 8 units of pseudo-rapidity, η. There are four subdivisions along the shower development axis in the electromagnetic (EM) section, and four or five in the hadronic (H) section. The hadronic calorimeter is divided into the fine hadronic (FH) section with relatively thin uranium absorber, and the backing coarse (CH) section. In the intercryostat region $0.8 < |\eta| < 1.6$ where the relatively thick cryostat walls give extra material for shower development, a scintillator based intercryostat detector (ICD) and extra 'massless gap' (MG) liquid argon gaps without associated absorber are located.

The calorimeter tower segmentation in $\eta x \phi$ is 0.1 x 0.1, which results in towers whose transverse size is larger than the expected sizes of EM showers but, considerably smaller than typical sizes of jets.

As a compromise, for triggering purposes, we add four adjacent calorimeter towers to form trigger towers (TT) with a segmentation of 0.2 x 0.2 in $\eta x \phi$. This yields an array that is 40 in η and 32 in ϕ or a total of 1,280 EM and 1,280 H tower energies as inputs to the L1 calorimeter trigger.

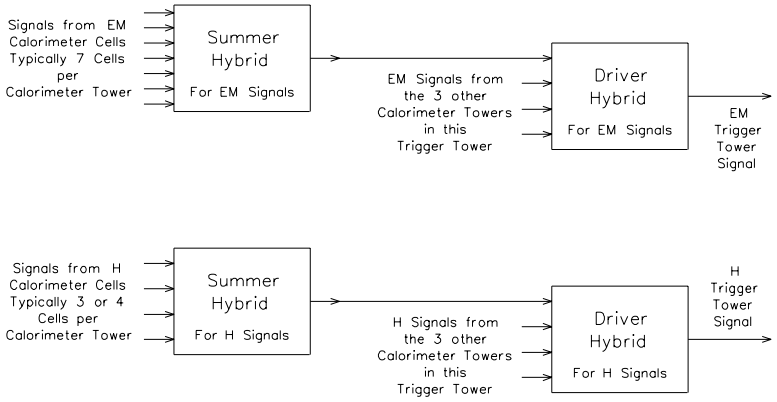

Figure 23. Trigger tower formation.

The analog summation of the signals from the various calorimeter cells in a trigger tower into the EM and H TT signals takes place as described on page 320. This arrangement for summing the calorimeter cells into trigger towers is shown schematically in Figure 23.

Long ribbons of coaxial cable route the 1280 EM and H analog trigger tower signals from the detector platform through the shield wall and then into the first floor of the moving counting house (MCH) where the Level 1 calorimeter trigger is located. The first step in the Level 1 calorimeter trigger is to scale these signals to represent the E_T of the energy deposited in each trigger tower and then to digitize these signals at the beam-crossing rate (132ns) with fast analog to digital converters. The digital output of these 2560 converters is used by the subsequent trigger logic to form the Level 1 calorimeter trigger decision for each beam crossing. The converter outputs are also buffered and made available for readout to both the Level 2 Trigger system and the Level 3 Trigger DAQ system.

The digital logic used in the Level 1 Calorimeter Trigger is arranged in a "pipe-lined" design. Each step in the pipe-line is completed at the beam crossing rate and the length of the pipe-line is less than the maximum DØ Level 1 trigger latency for Run IIa which is 3.3 μsec (driven by the calorimeter shaping times, cables lengths, drift times etc). This digital logic is used to calculate a number of quantities that are useful in triggering on specific physics processes. Among these are quantities such as the total transverse energy and the missing transverse energy, which we will designate as "global" and information relating to

"local" or cluster aspects of the energy deposits in the calorimeter. The latter would include the number of EM and H-like clusters exceeding a set of programmable thresholds.

4.3.2 Global Triggers

Interesting global quantities include:

the total transverse energies:

$$\left(E_T^{EM}\right)_{Total} = \sum_{i=1}^{1280} \left(E_T^{EM}\right)_i$$

$$\left(E_T^{H}\right)_{Total} = \sum_{i=1}^{1280} \left(E_T^{H}\right)_i$$

and

$$\left(E_T\right)_{Total} = \left(E_T^{EM}\right)_{Total} + \left(E_T^{H}\right)_{Total}$$

the missing transverse energy:

$$Mp_T = \sqrt{(E_x^2 + E_y^2)}$$

where:

$$E_x = \sum_{i=1}^{1280} \left[\left(E_T^{EM}\right)_i + \left(E_T^{H}\right)_i\right]\cos(\phi_i)$$

and

$$E_y = \sum_{i=1}^{1280} \left[\left(E_T^{EM}\right)_i + \left(E_T^{H}\right)_i\right]\sin(\phi_i)$$

Any of these global quantities can be used in constructing triggers. Each quantity is compared to a number of thresholds and the result of these comparisons is passed to the Trigger Framework where up to 128 different Level 1 triggers can be formed.

4.3.3 Cluster Triggers

The DØ detector was designed with the intent of optimizing the detection of leptons, quarks and gluons. Electrons and photons will manifest themselves as localized EM energy deposits and the quarks and gluons as hadron-like clusters.

Energy deposited in a Trigger tower is called EM-like if it exceeds one of the EM E_T thresholds and if it is not vetoed by the H energy behind it. Up to four EM E_T thresholds and their associated H veto thresholds may be programmed for each of the 1280 trigger towers. Hadronic energy deposits are detected by calculating the EM E_T + H E_T of each Trigger tower and comparing each of these 1280 sums to four programmable thresholds.

The number of Trigger towers exceeding each of the four EM thresholds (and not vetoed by the H energy behind it) is calculated and these four counts are compared to a number of count thresholds. The same is done for the four EM E_T + H E_T thresholds. The results of these count comparisons on the number of Trigger towers over each threshold are sent to the Trigger Framework where they are used to construct the Level 1 Triggers.

4.3.4 Hardware Implementation

4.3.4.1 Front End Cards

The analog signals from the calorimeter, representing energies, arrive at the Calorimeter Trigger over coaxial differential signal cables and are connected to the analog front end section of a Calorimeter Trigger Front End Card (CTFE). A schematic diagram of one of the four cells of this card is shown in Figure 24.

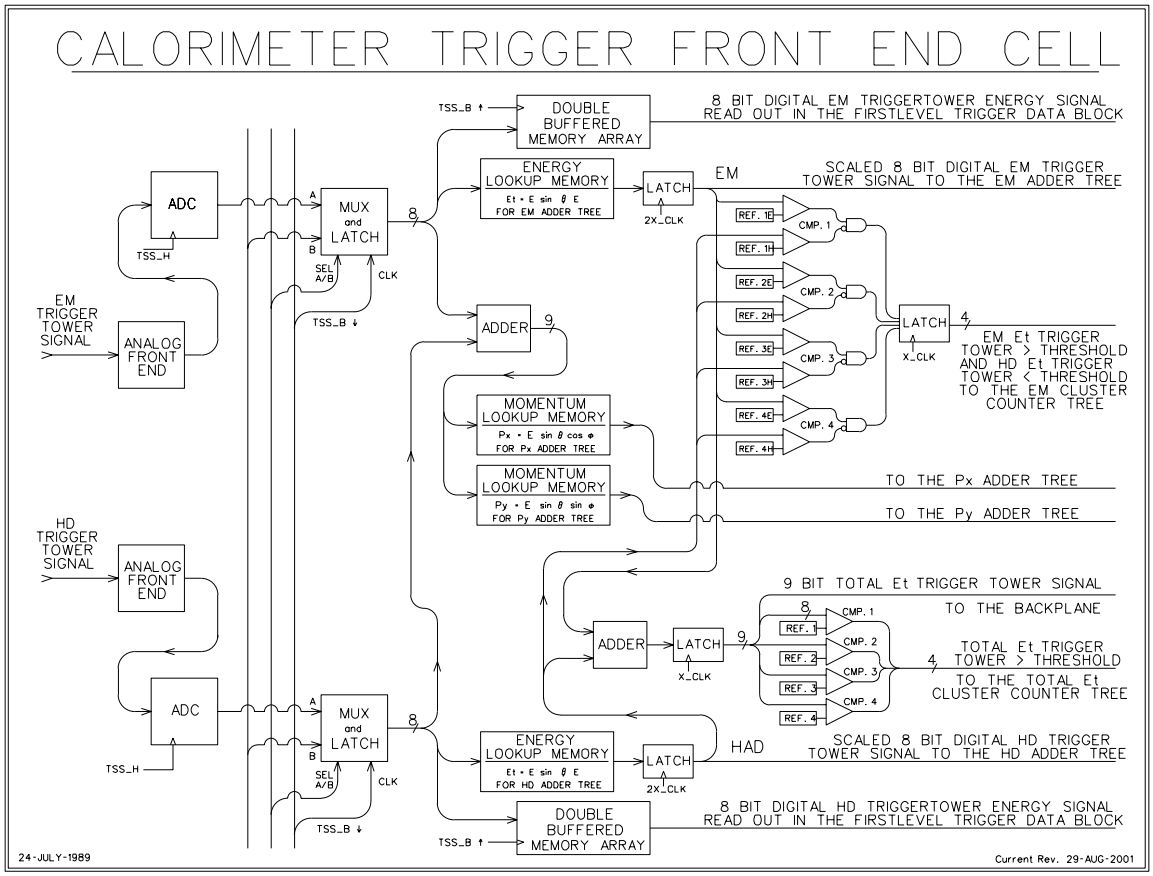

Figure 24. Calorimeter trigger front end cell (CTFE).

The front-end section contains a differential line receiver and scales the energy signal to its transverse component using a programmable gain stage. The front end also contains digital to analog circuitry for adding a positive bias to the tower energies in accord with downloaded values.

Immediately after the analog front end, the EM or H signal is turned into an 8 bit number by fast (20 ns from input to output) FADC's. With our current choice of 0.25 GeV least count this gives a maximum of 64 GeV for the single tower transverse energy contribution.

The data are synchronized at this point by being clocked into latches and then follow three distinct parallel paths. One of these paths leads to a pipeline register for digital storage to await the L1 trigger decision and subsequent readout to the Level 2 Trigger system and the Level 3 Trigger DAQ system.

On the other two paths, each 8-bit signal becomes the address to a look up memory. The content of the memory at a specified address in one case is the transverse energy with all necessary corrections such as lower energy requirements etc. In the other case, the EM + H transverse energies are first added and then subjected to two look-ups to return the two Cartesian components of the transverse energy for use in constructing MP_T. The inherent flexibility of this scheme has a number of advantages: any energy dependent quantity can be generated, individual channels can be corrected or turned off at this level and arbitrary individual tower efficiencies can be accommodated.

The CTFE card performs the function of adding the E_T's of the four individual cells for both the EM and H sections and passing the resulting sums onto the Adder Trees. In addition it tests each of the EM and EM+H tower transverse energies against the four discrete thresholds and increments the appropriate counts. These counts are passed onto the EM cluster counter trees and the total E_T counter trees, respectively.

4.3.4.2 Adder and Counter Trees

The adder and counter trees are similar in that they both quickly add a large number of items to form one sum. At the end of each tree the sum is compared to a number of thresholds and the result this comparison is passed to the Trigger Framework. A typical adder tree is shown in Figure 25.

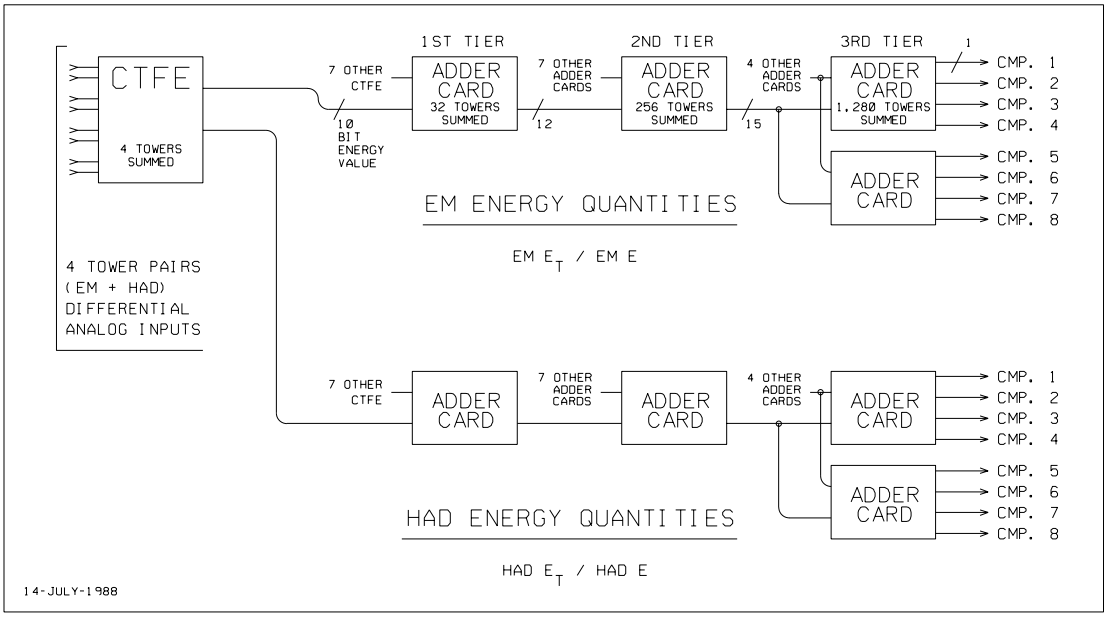

Figure 25. Adder tree for EM and Had quantities.

4.3.5 Physical Layout

Ten racks are used to hold the Level 1 Calorimeter Trigger, which is located in the first floor moving counting house. The lower section of each rack contains the CTFE cards for 128 Trigger towers (all 32 ϕ's for four consecutive η's). The upper section of each rack contains a component of one of the Adder or Counter Trees.

4.4 Motivations for Upgrading the Current System

The current L1 calorimeter trigger, which was built in 1988, and was used in Run 1 and Run IIa, has a number of features that limit its usefulness in Run IIb.

1) Trigger tower analog signals have rise times that are slightly longer than the 132 ns bunch spacing possible in Run IIb. The fall time of the signals, ~400 ns, is also significantly longer than the time between collisions. This makes it impossible for the current L1 calorimeter trigger to reliably assign calorimeter energy to the correct beam crossing, resulting in L1 trigger accepts being generated for the wrong beam crossing. Since information about the correct (interesting) beam crossing would be lost in these cases, finding a solution to this problem is imperative.

2) The fixed size trigger towers used in the current L1 calorimeter trigger are much smaller than the typical lateral size of a jet, resulting in extremely slow “turn-on” curves for jet and electron triggers. For example, a 6 GeV single tower threshold becomes ~100% efficient only for jets with transverse energies greater than 60 GeV. This poor resolution, convoluted with the steeply falling jet E_T spectrum, results in an overwhelming background of low energy jets passing a given threshold at high luminosity.

3) Total E_T and missing E_T resolution is significantly degraded because signals from the ICR detectors are not included in the trigger sums.

To run efficiently under all possible Run IIb conditions, the problem of triggering on the wrong bunch crossing must be resolved. Beyond that, the limited clustering capabilities in the current system result in unacceptably high rates for the triggers needed to discover the Higgs and pursue the rest of the DØ physics program. Each of these issues is discussed in more detail in the rest of this section, while our solutions are presented in the following sections.

4.4.1 Bunch Crossing mis-Identification

Because the width of the shaped analog TT signals is >400 ns, the current system will experience difficulties, as mentioned previously, if the spacing between bunches in the Tevatron is reduced from 396 ns to 132 ns. The main issue here is identifying energy deposited in the calorimeter with the correct bunch crossing. This is illustrated in Figure 21 and Figure 22, which show representative TT analog signals. The calorimeter readout timing is set such that the peak of the analog signal (~200 ns after it begins to rise) corresponds to the bunch crossing, *n*, where the relevant energy was deposited. For large amplitude signals, however, one or more of the TT E_T thresholds may be crossed early enough on the signal's rise to be associated with bunch crossing *n-1*. Additionally, the signal may not fall below threshold for several bunch crossings (*n+1, n+2,*…) after the signal peaks due to the long fall time. Because no events are accepted after an L1 trigger accept is issued until the silicon detector is read out, triggering on bunch crossing *n-1* would cause DØ to lose the interesting event at bunch crossing *n* in such a case.

4.4.2 Background Rates and Rejection

4.4.2.1 Simulation of the Current System

In order to assess the physics performance of the present L1 calorimeter trigger, the following simulation is used. The jet performance is studied using a Monte-Carlo sample of QCD events (PYTHIA, with parton p_T cuts of 5, 10, 20, 40 GeV and 0.5 overlaid minimum bias events). A cone algorithm with a radius of 0.4 in $\eta x \phi$ is applied to the generated stable hadrons in order to find the generated jets and their direction. The direction of each generated jet is extrapolated to the calorimeter surface; leading to the "center TT" hit by the jet. The highest E_T TT in a 3x3 trigger tower region (which is 0.6x0.6 in $\eta x \phi$ space) around this center is then used to define the "trigger E_T" corresponding to the jet.

4.4.2.2 Energy measurement and turn-on curves

In the present L1 calorimeter trigger, the trigger towers are constructed using fixed $\eta x \phi$ towers. Thus we expect that a trigger tower only captures a small fraction of the total jet energy since the size of the 0.2 x 0.2 trigger towers is small compared to the spatial extent of hadronic showers. This is illustrated in Figure 26, which shows, for simulated 40 GeV E_T jet events, the ratio of the E_T observed by the trigger to the generated E_T. It can be seen in Figure 26 that this transverse energy is only 25% of the jet E_T on average. Therefore we must use

low jet trigger thresholds if we are to be efficient even for relatively high energy jets. Moreover the trigger E_T has poor resolution, as can be seen in Figure 26. As a result, the trigger efficiency (the efficiency for having at least one TT with E_T above a given threshold) rises only slowly with increasing jet E_T, as shown in the turn-on curves in Figure 27. A similar effect occurs for the EM triggers as well; even though a typical EM shower can be reasonably well contained within a TT, often the impact point of an electron or photon is near a boundary between TTs.

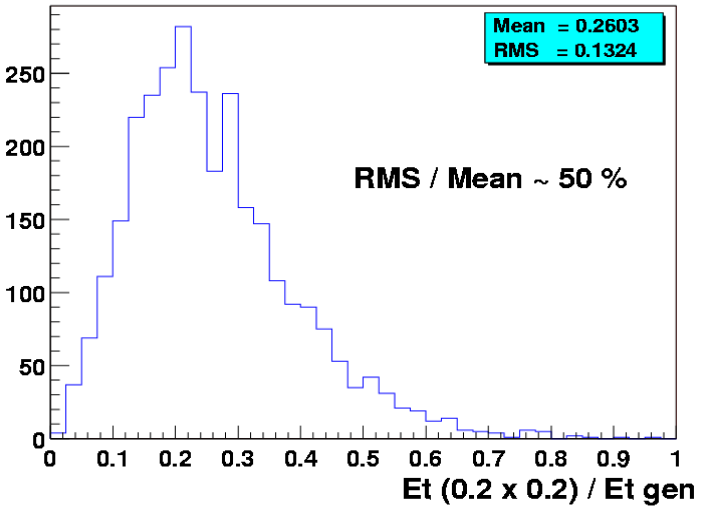

Figure 26. Ratio of the trigger E_T to the transverse energy of the generated jet. Only jets with $E_T \approx 40$ GeV are used in this figure.

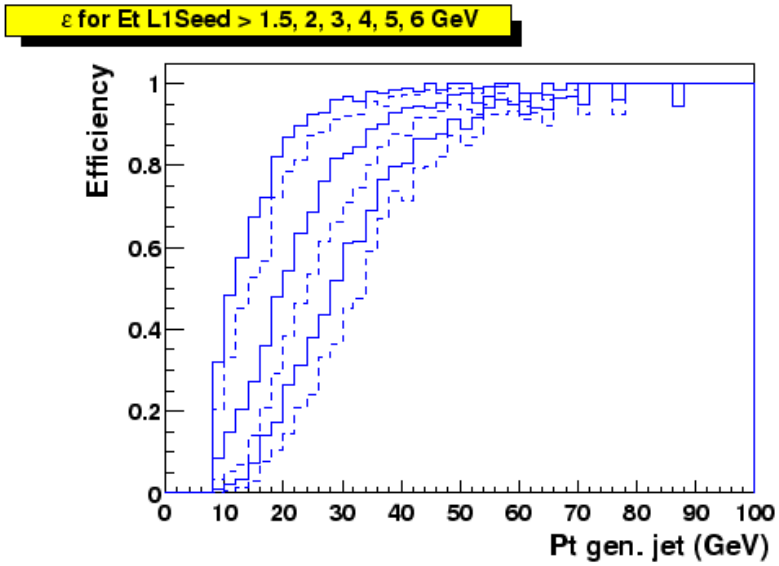

Figure 27. Trigger efficiency as a function of the transverse energy of the generated jet. The curves correspond to thresholds of 1.5, 2, 3, 4, 5 and 6 GeV (respectively from left to right).

4.4.2.3 Trigger rates

The trigger E_T resolution, convoluted with the steeply falling p_T spectrum of QCD events, leads to, on average, the "promotion" of events to larger E_T's than

the actual E_T. The number of QCD events which pass the L1 trigger is thus larger than what it would be with an ideal trigger E_T measurement. Due to the very large cross-section for QCD processes, this results in large trigger rates[4]. For example, as shown in Figure 28, an inclusive unprescaled high E_T jet trigger, requiring at least one TT above a threshold defined such that the efficiency for 40 GeV jets is 90%, would yield a rate for passing the L1 calorimeter trigger of at least 10 kHz at $2x10^{32}$ cm^2 s^{-1}. Maintaining this rate below 1 kHz would imply an efficiency on such high E_T jets of only 60%. Trigger rates increase faster than the luminosity due to the increasing mean number of interactions per bunch crossing. Trigger rates are shown in Figure 29 as a function of the mean number of minimum bias events which pile up on the high p_T interaction. These are shown for two multi-jet triggers: the first requiring at least two TT above 5 GeV (indicated as CJT(2,5)); the second requiring at least two TT above 5 GeV and at least one TT above 7 GeV (indicated as CJT(1,7)*CJT(2,5)). These triggers correspond to reasonable requirements for high p_T jets because, as can be seen in Figure 28, a threshold of 5 GeV leads, for 40 GeV jets, to an 80 % efficiency. The rates in Figure 29 are shown for a luminosity of $2x10^{32}$ cm^{-2} s^{-1}. For the higher luminosity of $5x10^{32}$ cm^2 s^{-1} expected in Run IIb, the L1 bandwidth of 5kHz could be saturated by such dijet conditions alone, unless large prescale factors are applied.

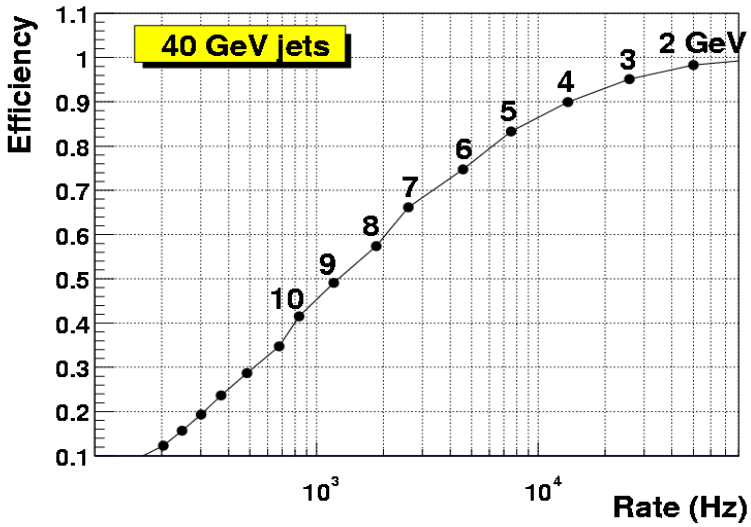

Figure 28. The efficiency to trigger on 40 GeV jets as a function of the inclusive trigger rate when one TT above a given threshold is required. Each dot corresponds to a different threshold (in steps of 1 GeV), as indicated. The luminosity is $2x10^{32}$ $cm^{-2}s^{-1}$.

[4] These rates are estimated here from samples of PYTHIA QCD events with parton p_T > 2GeV, passed through a simulation of the trigger response.

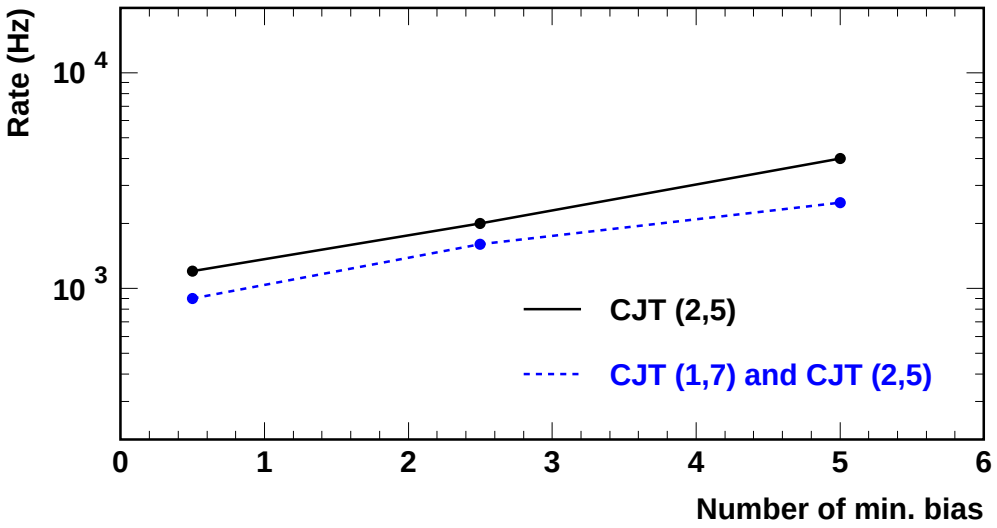

Figure 29. The inclusive trigger rate as a function of the mean number of minimum bias events overlaid on the high p_T interaction. The rates are shown for two di-jet trigger conditions corresponding to two TTs above 5 GeV (CJT(2,5)) and two TTs with above 5GeV and at least one above 7 GeV (CJT(1,7)*CJT(2,5)). The luminosity is $2x10^{32}$ cm^{-2} s^{-1}.

A more exhaustive study of the evolution of the L1 trigger rate with increasing luminosity has been carried out[5]. In that document a possible trigger menu was considered, in which ~75 % of the L1 bandwidth is used by multijet triggers. The results are shown in Table 5. It can be seen that, at the luminosity foreseen for Run IIb (corresponding to the 4th row), the trigger rates should be reduced by at least a factor of four in order to maintain a reasonably small dead time. We note that the need to preserve jet triggers is required by some of the Higgs boson physics (for example, $p\bar{p} \to ZH \to \nu\bar{\nu}b\bar{b}$).

Table 5. The overall level 1 trigger rates as a function of luminosity.

Luminosity	High Pt L1 rate (Hz)	Total L1 rate (Hz)
$1x10^{32}$ cm^{-2} s^{-1}.	1,700	5,000
$2x10^{32}$ cm^{-2} s^{-1}.	4,300	9,500
$5x10^{32}$ cm^{-2} s^{-1}.	6,500	20,000

4.4.3 Conclusions/implications for high luminosity

Clearly, the bunch crossing mis-identification problem must be resolved for Run IIb or the L1 calorimeter trigger will cease to be effective. The physics studies presented above also show that there is a need to significantly improve the rejection of the L1 calorimeter trigger (while maintaining good efficiency) if we

[5] B. Bhattacharjee, "Transverse energy and cone size dependence of the inclusive jet cross section at center of mass energy of 1.8 TeV", PhD Thesis, Delhi University.

are to access the physics of Run IIb. One obvious way to help achieve this is to migrate the tools used at L2 (from Run IIa) into L1. In particular, the ability to trigger on "objects" such as electromagnetic showers and jets would help significantly. The "clustering" of TT's at L1, could reduce the trigger rates by a factor 2 to 4 as will be shown later. The principal reason for this gain comes from the improvement in the quality of the energy cut, when applied to a cluster of trigger towers. Transferring to level 1 some of the functions that currently belong to level 2 would also permit the introduction of new selection algorithms at the L1 trigger level. So while it is clear that there are additional gains to be made through EM trigger tower shape cuts and missing E_T filtering, they will require further study to quantify the specific gains. These studies remain to be done.

From a conceptual viewpoint, an important consequence of selecting physics "objects" at level 1 is that it allows a more "inclusive" and hence less biased selection of signatures for the more complicated decays to be studied in Run IIb. Thus we expect that the trigger menus will become simpler and, above all, less sensitive to biases arising from the combinations of primary objects.

4.4.4 Overview of Calorimeter Trigger Upgrade

We have examined various possibilities for the changes necessary to address the incorrect bunch crossing assignment problem at 132 ns bunch spacing and the trigger energy resolution problem. The age and architecture of the current system prohibit an incremental solution to these issues. Changing one aspect of the system, for example implementing a new clustering algorithm, has an impact on all other parts since in the current electronics both digitization of TT signals and the search for TTs above threshold happens on the same board. We therefore propose to design and build an entirely new L1 calorimeter trigger system, which will replace all elements of the current trigger downstream of the BLS cables. A partial list of improvements provided by the new system is given below.

- Necessary hardware improvements in filtering to allow proper triggering on the correct bunch crossing at 132 ns bunch spacing.
- Implementation of a "sliding window" algorithm for jets, electrons and taus.
- The addition of presently unused calorimeter energy information from the intercryostat detector (ICD) and massless gaps (MG) in the L1 trigger.
- Optimization of trigger tower thresholds.
- The ability to better correlate tracks from the fiber tracker to calorimeter clusters.

Studies of these improvements are discussed in the following sections with the exception of the correlation between tracks and calorimeter clusters, which is described in the Cal-Track Match system section. The design of the new system is then outlined in the remaining parts of the chapter.

4.5 Digital Filtering

Digital filtering offers a way to reduce the effect of unwanted triggers due to collisions in close proximity to the desired trigger.

4.5.1 Concept & physics implications

The pulse shape, and particularly the rise time, of the trigger pickoff signal is not optimized for 132ns beam bunch crossing operation (see Figure 21 and Figure 22). Since the trigger pickoff pulse width significantly exceeds the 132ns bunch spacing time of Run IIb, the ability to correctly identify the correct trigger bunch crossing is compromised. There may be intermediate solutions to address this problem at the lower luminosities, but a long-term solution must be developed. This could be done by means of an analog filter with shorter shaping, but this is only achieved with a further loss in signal. A digital filter is a better solution because it is much more flexible for a similar cost.

The trigger pickoff signal is at the end of the calorimeter electronic chain described above. The ideal energy deposition shape is a "saw-tooth" pulse (infinitely fast rise and a linear ~400ns fall) from energy deposited in the cells of the calorimeter at each beam crossing. This is modified by the transfer function of the electronics. The inverse transfer function will transform the pickoff signal back to original energy deposition pulse shape. Digital filtering would be implemented at this stage. The inverse function can be implemented by a FIR (Finite Impulse Response) digital filter. In the presence of noise, the digital filter offers an additional advantage: one can use the theory of optimal filtering to minimize the noise contribution.

In order to define the exact form of a digital filter best suited to the task, a measurement of noise in the trigger pickoff signals is needed. As such measurements become available, a refined design will be undertaken.

4.5.2 Pileup rejection

Two different "pile-up" effects arise with increasing luminosity, the first is due to extra collisions in the crossing of interest (and thus unavoidable), and the second is due to collisions in neighboring crossings that contribute to the crossing of interest because of signal shapes.

In the first case, we find that as the luminosity increases, then for each triggered beam crossing there are several minimum bias events that appear in that same beam crossing. The number of such additional events is Poisson distributed with a mean proportional to the luminosity. The energy added by these events has a distribution close to that of a double exponential (Laplacian). It is possible to minimize the contribution of this noise by using an appropriate digital filter (Matched Median Filter).

In the second case, because the width of the trigger pickoff signal extends over several beam crossing (6 at 132ns on the positive side of the signal), then when two such pulses are close in time, there is some overlap and thus the shape of the pickoff signal becomes more complicated than that of a single isolated pulse. The inverse filter will extract from this signal the two original

pulses. Consequently, the problems caused by overlapping pulses are minimized if one uses digital filtering.

4.5.3 Input data and simulation tools

A series of measurements on several EM and HAD trigger pickoff channels was performed to provide the necessary input to digital filter algorithm studies. Oscilloscope traces and raw data files have been recorded. A chain of programs has been developed to generate training sets based on measured pulses, simulate the analog to digital conversion stage, study digital filter algorithms and compare results with the expected outputs. All programs are standalone and use ASCII files for input and output to provide an increased flexibility and the widest choice of tools for visualization and post-processing.

A typical pulse on an EM channel is shown on the left side of Figure 30. A 4096-point Fast Fourier Transform of this signal is shown on the right side of Figure 30 (the DC component was removed for clarity). It can be seen that most of the energy of the signal is located in frequency components below ~10 MHz. In order to remove the high frequency noise that can be seen, we suggest that an analog low-pass filter is placed on each channel before the analog to digital converter. Different filters were investigated by numerical simulation. As shown on the figure, a 2nd order low-pass Butterworth filter with a cutoff frequency of 7.57 MHz seems adequate to remove high frequency oscillations on the signal while preserving the shape of its envelope. Such low-pass filter will avoid the potential problems of spectrum aliasing in the digital domain.

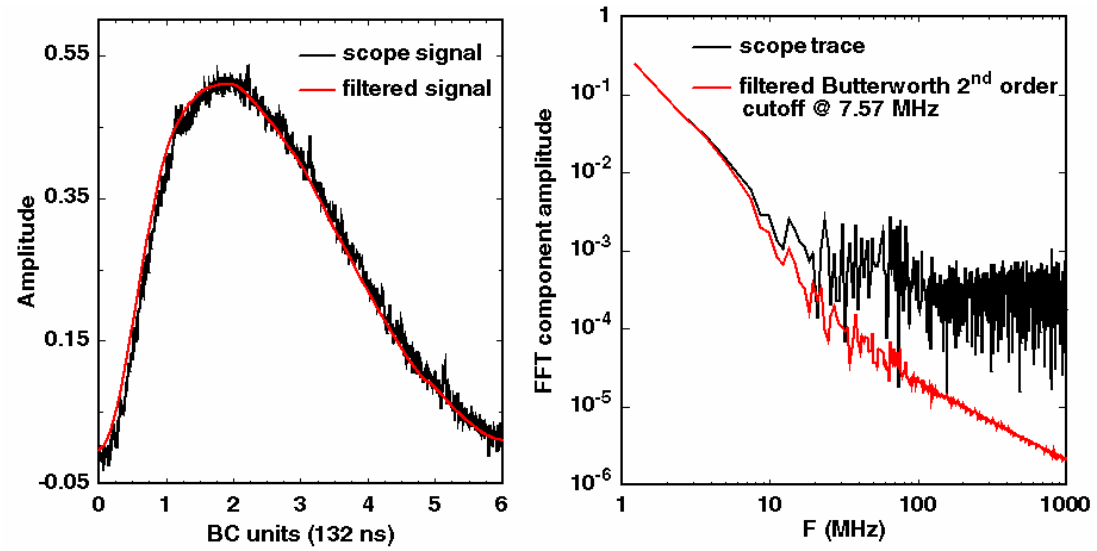

Figure 30. Scope trace of a typical EM pulse and corresponding spectrum. Pulse shape and spectrum after an anti-aliasing filter.

4.5.4 Algorithm evaluation parameters

In order to investigate and compare digital filter algorithms, several criteria have been defined. A first set is related to the features of the algorithm itself:

irreducible latency, number of parameters to adjust and channel dependency, procedure for parameter determination and tuning, operating frequency, behavior under digital and analog saturation… A second family of criteria relates to the quality of the algorithm: precision on the estimated Et value for the beam-crossing of interest and residual error on adjacent beam-crossings, time/amplitude resolution, ability to separate pulses close in time, probability of having pulses undetected or assigned to the wrong beam-crossing. Several criteria are related to the sensitivity of an algorithm: robustness against electrical noise, ability to reject pileup noise, sensitivity to signal phase and jitter with respect to a reference clock, dependence on pulse shape distortion, performance with limited precision arithmetic, influence of coefficient truncation and input quantization, etc. The last set of comparison criteria concerns implementation: amount of logic required and operating speed of the various components, effective latency.

Defining and selecting the algorithm that will lead to the best trade-off between all these – sometimes contradictory – criteria is not straightforward. Some compromises on performance and functionality will necessarily be done in order to fit in the tight, non-extensible, latency budget that can be devoted to this task while keeping the system simple enough to be implemented with modern, industrial electronic devices at an affordable cost. Algorithm definition and test by computer simulation, electronic hardware simulation and validation with a prototype card connected to real detector signals are among the necessary steps for a successful definition of the digital filter.

4.5.5 Algorithms studied

At present, three types of algorithms have been proposed and investigated. These are:

- A Finite Impulse Response (FIR) deconvolution filter;
- A peak detector followed by a weighed moving average filter;
- A matched filter followed by a peak detector.

We describe these algorithms and their simulation studies of their performance below. Based on these studies, the matched filter algorithm has been selected for the baseline calorimeter trigger design.

4.5.5.1 FIR deconvolution

The deconvolution filter is designed to implement the inverse transfer function of the complete calorimeter pickoff chain. When driven with a typical trigger pickoff saw-tooth shaped pulse, the output of the filter is the original pulse. In order to produce a meaningful output for each beam crossing, the filter must have a bandwidth equal at least to the beam crossing frequency. Hence, input samples must be acquired at least at twice the beam-crossing rate (Shannon's sampling theorem). However, only one output value per beam crossing is computed. Coefficient count must be sufficient to ensure that the output of the filter remains null during the falling edge of the input pickoff signal. The determination of coefficients can be made using a set of input training samples

that include noise, pileup, time jitter and pulse shape distortion. The differences between the expected values and the actual filter outputs are accumulated and a least mean square minimization is performed to determine the set of coefficients that provide the optimum solution.

The deconvolution filter is linear; parameter tuning can lead to the optimum linear solution for the set of input constraints that is given. This filter performs well to separate pulses close in time as illustrated in Figure 31. A series of pulses of constant height separated by a decreasing amount of time were generated and a simulated trigger pickoff signal was calculated. It can be seen in Figure 31 that the deconvolution FIR filter is able to identify correctly adjacent pulses, even when these occur on two consecutive beam-crossings (i.e. 132 ns apart). However, a non-null residual error is present for some beam-crossings.

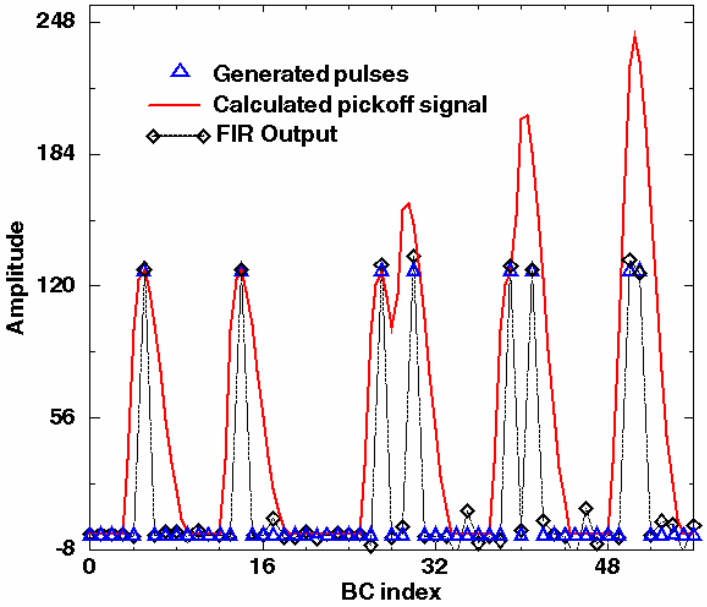

Figure 31. Deconvolution of pulses overlapping in time. Sampling rate is 15.14 MHz (BC x 2), ADC precision is 8 bit; a 12-tap FIR is used; 32-bit floating-point arithmetic is used for coefficients and computations.

Various tests were performed to investigate the behavior and the performance of the FIR deconvolution algorithm. An example is shown in Figure 32. In this test, filter coefficients are optimized for a given pulse shape (no noise and no time jitter in the training set), with the peak of the signal precisely phased-aligned with the analog to digital converter sampling clock. A train of pulses of constant amplitude (128 on an 8-bit range) with a phase varying in [-1/2 BC, ½ BC] with respect to the sampling clock is generated. Two sets of observations are distinguished: the value of the output for the beam-crossings that correspond to a simulated deposition of energy and the residual error for the beam-crossings where a null response is expected. For a null phase, it can be seen in Figure 32 that the output of the filter corresponds to the expected output for the beam-crossing of interest and is null for adjacent beam-crossings. When the phase is varied, not only a growing error is made on the energy estimated for the correct

BC, but also a non-null output for adjacent BC's is observed. The algorithm is somewhat sensitive to sampling clock phase adjustment and signal jitter. A possible improvement would be optimize the filter coefficients with a training set of samples that include time jittered pulses.

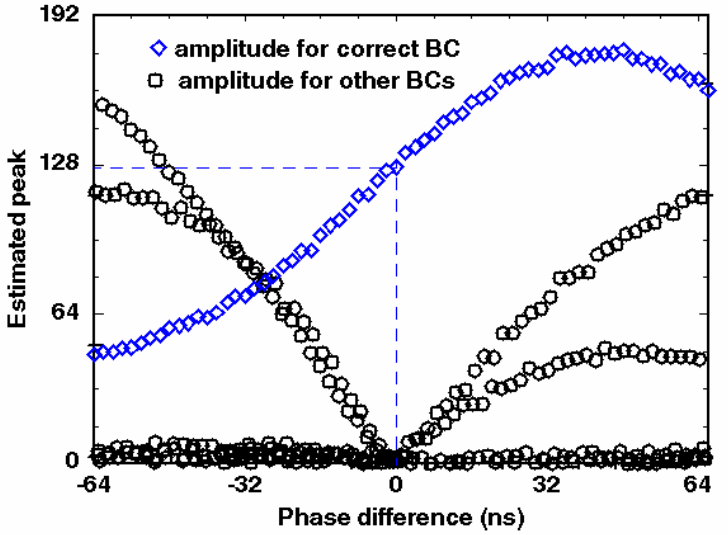

Figure 32. Operation of a deconvolution FIR filter when the phase of pulses is varied. Sampling rate is 15.14 MHz (BC x 2), ADC precision is 8 bit; a 12-tap FIR is used; 32-bit floating-point arithmetic is used for coefficients (signed) and computations.

Other difficulties with the deconvolution FIR filter include its sensitivity to limited precision arithmetic and coefficient truncation, degraded performance when signal baseline is shifted and misbehavior when the input is saturated. Implementation is also a potential issue because a filter comprising over 12-tap is needed (assuming input samples are acquired at BC x 2). Although only one convolution product needs to be calculated per BC, a significant amount of resources would be needed to compute the corresponding 7.57 x 12 = 90 million multiply-accumulate operations per second per channel. Although an independent study of this algorithm could bring better results, linear deconvolution is not seen as a satisfactory solution.

4.5.5.2 Peak detector + weighed moving average

This non-linear filter comprises two steps: detecting the presence of a peak and calculating its height. The first step is accomplished by comparing the magnitude of ~3-4 successive samples. There is no specific method to pick up the minimum sets of conditions that these amplitudes need to satisfy to characterize the presence of a peak. Let E(kT) be the amplitude of input sample k. A possible set of conditions for peak detection can be expressed as follows:

A peak is present at t=(k-1)T IF

E(kT) < E[(k-1) T] AND

E[(k-1) T] >= E[(k-2) T] AND

E[(k-2) T] >= E[(k-3) T]

This set of conditions determines the presence of a peak with an irreducible latency of one period T. Empirical studies were performed to determine a viable set of conditions and a satisfactory sampling period T. The conditions mentioned above were retained; sampling at BC x 3 was chosen.

The second part of the algorithm consists in assigning 0 to the output if the peak detector did not detect the presence of a peak, or calculate the weighed average of several samples around the presumed peak. To simplify computations, the sum can be made over 4 samples with an identical weight of ¼ for each of them. A common multiplicative scaling factor is then applied.

One of the tests performed with this algorithm is shown in Figure 33. A series of pulses of growing amplitudes is generated. It can be seen that small pulses are not well detected. It should be also observed that, as expected, the output is null between pulses.

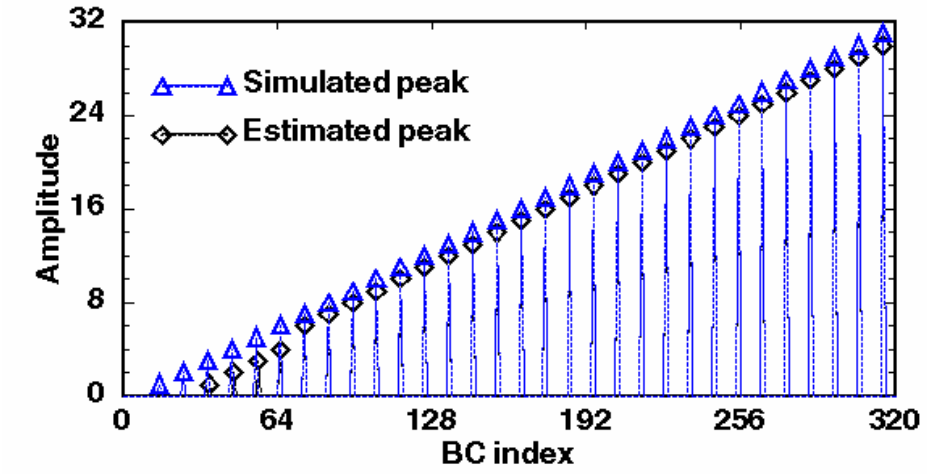

Figure 33. Operation of a peak detector + moving average. Sampling rate is 22.71 MHz (BC x 3), ADC precision is 8 bit; average is made on 4 samples; each weight is ¼; a common 8-bit multiplicative factor is applied; fixed-point arithmetic is used.

Other tests show that this algorithm is rather tolerant to signal phase and jitter, does not depend too much on pulse shape (except for double peaked HAD channels), is very simple to implement and has a low latency. Its main limitations are the empirical way for parameter tuning, low performance for small signals, misbehavior in case of pileup, the assignment of energy to the beam-crossing preceding or following the one of interest in some cases, the possibility that a pulse is undetected in some other cases. Although this algorithm is acceptable in some cases, it does not seem sufficiently robust and efficient.

4.5.5.3 Matched filter + peak detector

This algorithm comprises two steps. The matched filter is designed to best optimize the Signal to Noise Ratio when detecting a signal of a known shape degraded by white noise. In this case, it can be shown that the optimal filter for a signal E(kT) is the filter whose impulse response is:

h(kT) = E(T0 – kT) where T0 is a multiple of the sampling period T and is selected to cover a sufficient part of the signal. Because T0 has a direct influence on the irreducible latency of the algorithm, the number of filter taps and the operating frequency of the corresponding hardware, its value should be carefully chosen. The parameters to determine are: the sampling period T, the number of samples during T0, and the phase with respect to the training pulse of the temporal window used to determine the impulse response of the matched filter. It should also be mentioned that the peak produced at the output of a matched filter occurs at (nT+T0) and that this irreducible latency does not correspond to a fixed delay with respect to the occurrence of the peak in the signal being detected when some of the parameters of the filter are changed. When designing a series of filters running in parallel, care must be taken to ensure that algorithm latency is identical for all channels.

The second step of the algorithm is peak detection. A possible algorithm is the 3-point peak detector described by the following pseudo-code:

Peak present at t=(k-1)T IF E(kT) < E[(k-1) T] AND E[(k-1) T] > E[(k-2) T]

This peak-detector adds one period T of irreducible latency. If the conditions that characterize a peak are not satisfied, the output is set to 0, otherwise it is assigned the value of the matched filter.

Figure 34 where pulses of growing amplitudes (up to $1/8^{th}$ of the full 8-bit scale) have been generated. It can be seen that the algorithm performs well in that range. All pulses but the smallest ones have been correctly detected and have been assigned to the correct beam crossing. The output is exactly zero for the beam-crossings around those of interest. Intuitively, one can easily understand that the capability to produce minimal width pulses (one sample width) surrounded by strictly null outputs is more easily achieved with a non-linear filter than with a linear algorithm.

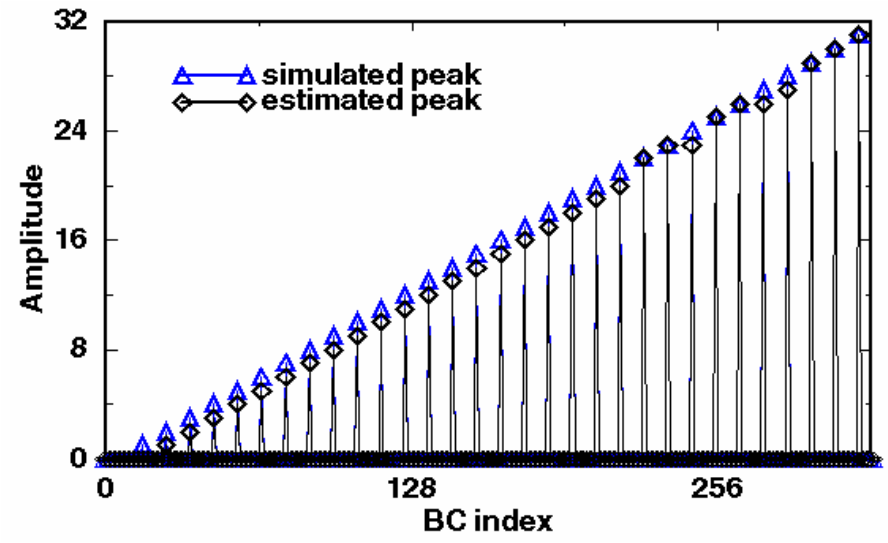

Figure 34. Operation of a matched filter + peak detector. Sampling rate is 15.14 MHz (BC x 2), ADC precision is 8 bit; 6 6-bit unsigned coefficients are used; fixed-point arithmetic is used.

The sensitivity of the matched filter and peak detector to signal phase shift was studied. Pulses of constant amplitude (1/2 full scale) and variable phase were generated. The relative error on the reconstructed amplitude for the beam crossing of interest is plotted in Figure 35. It can be seen that the relative error is confined within 5% when the phase shift is in the interval [-32 ns, 32 ns]. For larger phase shift values, the pulse is undetected and the output of the filter is null. This is a case of severe failure for this algorithm. For the beam-crossings surrounding that of interest, the output of the filter remains null over the range of phase shifts simulated; no erroneous assignment to the preceding or following beam crossing were observed.

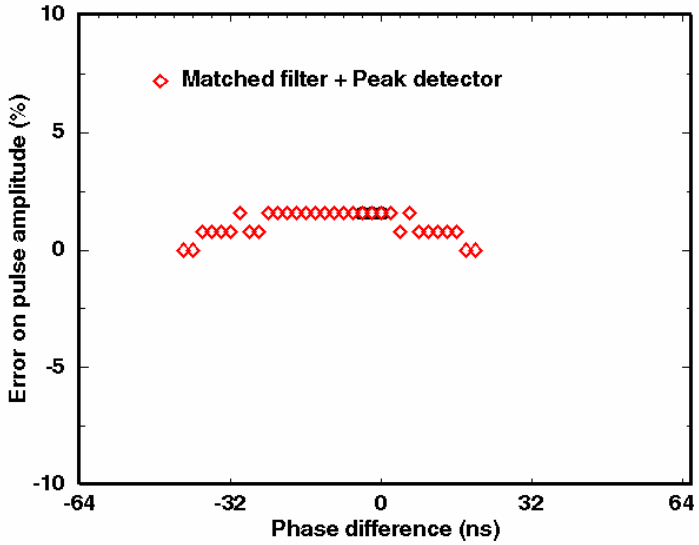

Figure 35. Operation of a matched filter + peak detector when signal phase is varied. Sampling rate is 15.14 MHz (BC x 2), ADC precision is 8 bit; a 6-tap matched filter with 6-bit unsigned coefficients followed by a 3-point peak detector are used; all computations are done in fixed-point arithmetic.

By comparing these results with that of the FIR deconvolution shown in Figure 32 (where the absolute value of filter output is plotted), it can be concluded that the matched filter algorithm is much more tolerant to signal phase and jitter.

Determining the number of taps for the matched filter requires making a compromise between the quality of the results, the latency of the algorithm and the amount of resources needed for implementation. A test was made to investigate the influence of the number of filter taps. A series of pulses of growing amplitudes (full 8-bit range) were generated. The reconstructed amplitude is shown in Figure 36 for a matched filter with 8-taps and 5-taps respectively. No significant degradation of performance was observed as long as the number of coefficients is greater or equal to 5. The difference in latency between the 8-tap version and the 5-tap version is 1 BC; the amount of computation to perform is increased by 60% when the number of taps is changed from 5 to 8.

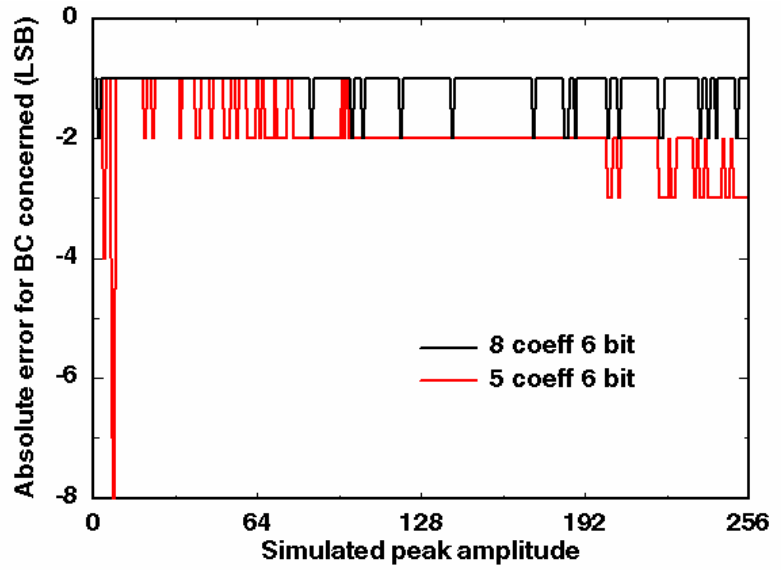

Figure 36. Operation of a matched filter + peak detector with different number of taps. Sampling rate is 15.14 MHz (BC x 2), ADC precision is 8 bit; coefficients are 6-bit unsigned; fixed-point arithmetic is used.

Algorithm behavior in case of saturation is also an important parameter. A series of pulses with amplitude that goes up to twice the range of the ADC (8-bit in this test) was generated. A comparative plot for the 3 algorithms studied is shown in Figure 37. The FIR deconvolution filter has two annoying features: the amplitude estimated for the BC of interest decreases, and the estimation on adjacent BC's grows rapidly as the level of saturation is increased. The peak detector has a satisfactory behavior under moderate saturation, but the peak of energy is assigned to the wrong beam crossing when the saturation level is increased. The matched filter has a smoothly growing output, and still assigns the energy value to the correct beam-crossing under a high level of saturation. Although in real experimental conditions, the combined effects of analog and digital saturation will be much more complex than what was simulated, the matched filter clearly appears to be superior to the two other algorithms.

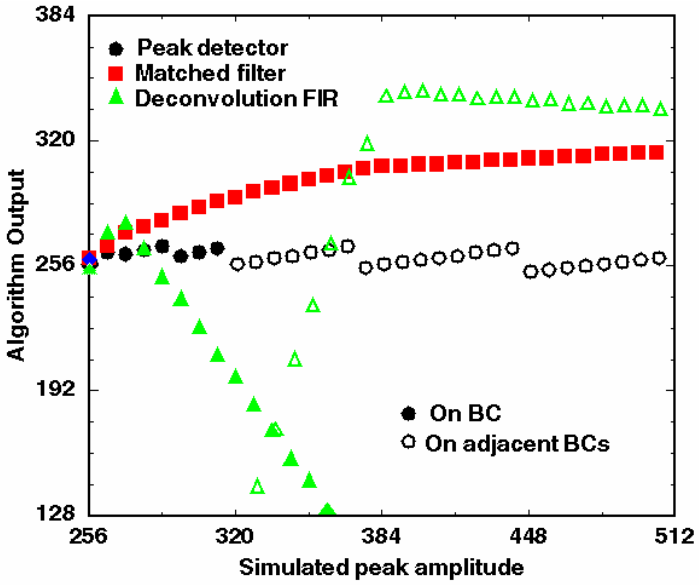

Figure 37. Algorithm behavior under digital input saturation.

Electronic noise reduction and pileup rejection are other properties that need to be considered to select the proper algorithm for digital filtering. At present, almost no studies have been made in these fields but a few simple tests. A series of couple of pulses of constant height (1/2 full range) separated in time by 10, 3, 2 and 1 beam-crossings have been generated. The output for the 3 algorithms studied is shown in Figure 38. As previously mentioned, the deconvolution FIR filter is able to correctly identify pulses that are close in time. On the other hand, both the peak detection scheme and the matched filter algorithm fail to identify the two pulses and their amplitude when pickoff signals overlap. One of the two pulses is systematically dropped and the energy of the remaining pulse is overestimated by a large factor. This configuration corresponds to a case of failure for these two algorithms. Detailed studies are needed to determine what will be the noise and pileup conditions in the real experiment and decide if the level of algorithm failures observed is below an acceptable limit or not. Tests with in the experiment with real signals are also crucial.

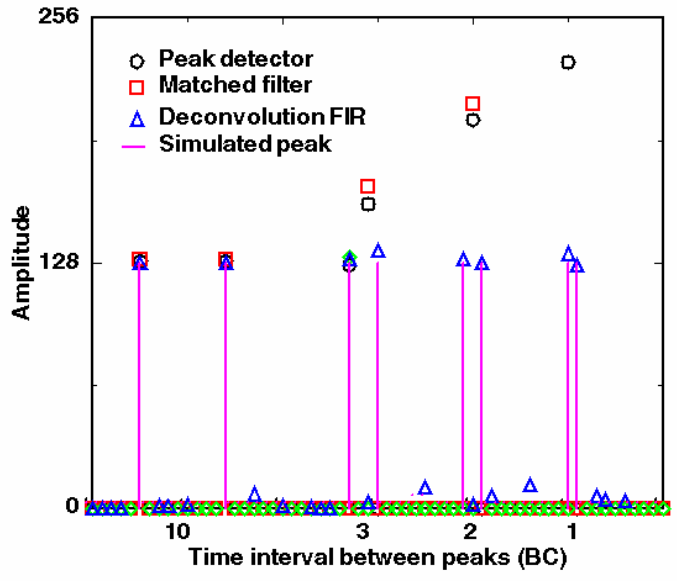

Figure 38.Behavior of the 3 algorithms with pulses close in time.

In order to compare the 3 algorithms studied, 8 criteria of merit were selected and subjective marks between 0 and 5 (0 is worse, 5 is best) were given for each algorithm. The resulting diagram is plotted in Figure 39. While none of the algorithm performs best in all fields, the optimum algorithm is the one whose polygon covers the largest area. Clearly, the matched filter is the algorithm that offers the best trade-off between all criteria. This algorithm is therefore the baseline for the prototype that is being designed and that will be tested in-situ.

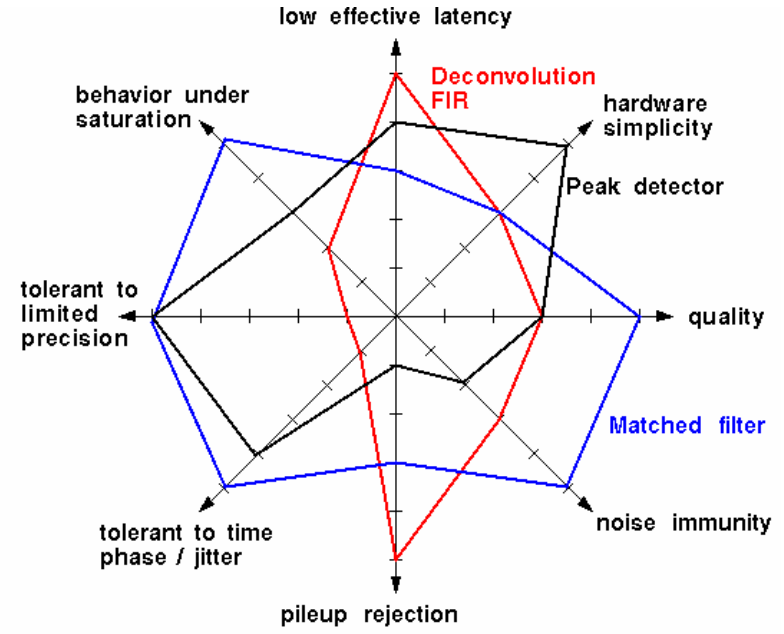

Figure 39. Comparison of the 3 algorithms proposed against 8 criteria of merit.

A number of studies still need to be done to confirm that this algorithm is the most appropriate. These include noise and pileup studies, the possibility to run computations at the beam-crossing rate (instead of running at BC x 2), the development of a scheme and program for the automatic determination of filter coefficients, etc. Tests on the detector must also be done and analyzed. While simulation offers a very good mean of investigation, real signals shapes, noise, time jitter, pileup conditions and many other effects cannot be taken into account without performing a real test. At present, studies on the digital filter allowed to select a candidate algorithm. In the prototype implementation, some flexibility will be allowed at that level, but algorithm changes will be confined to the capability of the available programmable logic.

4.5.6 Conclusions

Given the relatively slow trigger sum driver pulse shapes observed in Figure 16 and Figure 17, we believe that a digital filter is required to suppress the contributions from signals in nearby bunch crossings to that containing a high pT trigger. The matched filter algorithm offers the best performance among the digital filter algorithms studied and has been selected. Details of the implementation of the digital filter are given in Section 4.7.5.

4.6 Clustering algorithm simulation results

Algorithms relying on "sliding" trigger towers (TTs) can significantly improve the trigger performances, compared to the current calorimeter trigger based on single 0.2 x 0.2 TTs by better identifying the physical objects. Such algorithms have been extensively studied for the Atlas experiment, as described in the ATLAS Level-1 Trigger Technical Design Report[6].

Various algorithms can be used to cluster the trigger towers and look for "regions of interest" (***R***), i.e. for regions of fixed size, ***S***, in $\eta x \phi$ in which the deposited E_T has a local maximum. To find those regions of interest, a window of size ***S*** is shifted in both directions by steps of 0.2 in η and ϕ. By convention each window is unambiguously (although arbitrarily in the 2 x 2 case) anchored on one trigger tower T and is labeled S(T). Examples are shown in Figure 40.

[6] "The ATLAS Level-1 Trigger Technical Design Report", http://atlasinfo.cern.ch/Atlas/GROUPS/DAQTRIG/TDR/tdr.html, June 1998. See also "Trigger Performance Status Report", CERN/LHCC 98-1

Figure 40. Examples of (a) a 3x3 and (b) 2x2 sliding window S(T) associated to a trigger tower T. Each square represents a 0.2 x 0.2 trigger tower. The trigger tower T is shown as the shaded region.

The sliding windows algorithm aims to find the optimum region of the calorimeter for inclusion of energy from jets (or EM objects) by moving a window grid across the calorimeter η, ϕ space so as to maximize the transverse energy seen within the window. This is simply a local maximum finding algorithm on a grid, with the slight complication that local maxima are found by comparing energies in windows, S(T), which may have overlapping entries. Care must therefore be taken to ensure that only a single local maximum is found when two adjacent windows have the same S(T). A specific example of how the local maximum could be defined is shown in Figure 41. This process, which avoids multiple counting of jet (or EM object) candidates, is often referred to as "declustering".

The window corresponding to a local maximum is called the region of interest, ***R***, and is referenced by a specific TT within ***R*** as indicated in Figure 40 for a 3x3 or a 2x2 window. The total E_T within ***R*** plus that within a defined neighbor region is termed the trigger cluster E_T relevant to the jet or EM object.

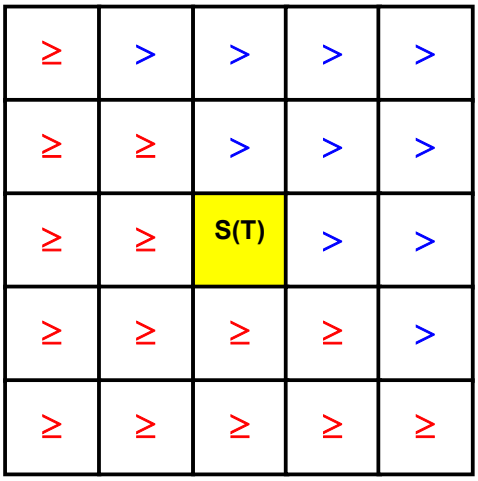

Figure 41. An illustration of a possible definition of a local E_T maximum for a R candidate. The cluster S(T) is accepted as an *R* candidate if it is more energetic than the neighboring clusters marked as ">" and at least as energetic as those marked "≥". This method resolves the ambiguities when two equal clusters are seen in the data. In this example, the declustering is said to be performed in a window of size 5x5 in $\eta x \phi$.

"Sliding window" algorithms can be labeled by three numbers – *x,y,z*, where:

- *x* = the size in $\eta x \phi$ of the sliding window's S(T). *x*=2 means that the sliding windows are defined by summing E_T's in 2×2 TTs in $\eta \times \phi$.
- *y* = the minimum overlap, in TTs, allowed between two sliding windows which can be considered to be regions of interest (local maxima). This is related to the size of the sliding window and the number of windows that are compared to determine if a given S(T) is a local maxima (see Figure 41). For sliding windows with *x*=2, *y*=1 means that local maxima must be separated by at least 2 TTs (one beyond the edge of the window). This corresponds to a declustering region of 5×5 windows.
- z = the size of the ring of neighbor TTs whose energy is added to that of ***R***, to define the trigger E_T. For sliding windows with *x*=2, *z*=1 means that trigger cluster E_T's are calculated over a 4×4 region.

Specific parameters which enter in the definition of the electromagnetic or tau triggers only will be detailed in the relevant sections below.

4.6.1 Jet algorithms

We have chosen to implement the "2,1,1" scheme as our baseline jet algorithm. This algorithm is shown schematically in Figure 42.

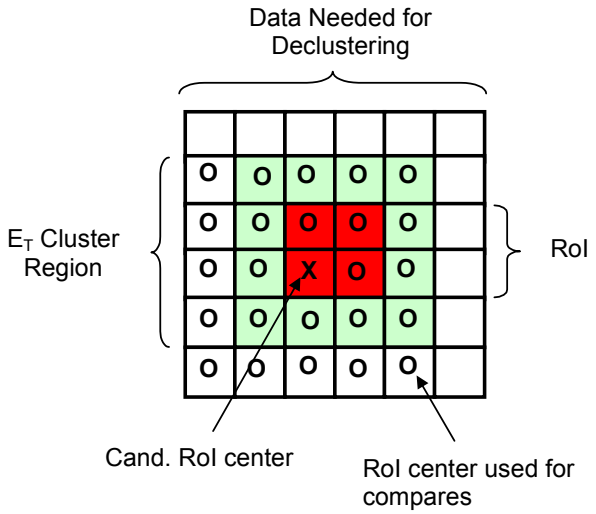

Figure 42: Schematic representation of the baseline jet algorithm.

The reasons for this choice are detailed in the following sections where several possible sliding windows algorithms are compared. The parameters of these algorithms are shown in Table 6.

Table 6: Details of some of the algorithms compared in the following plots.

Algorithm	Window Size ($\eta\times\phi$ TTs)	Declustering Region ($\eta\times\phi$windows)	Trigger Cluster ($\eta\times\phi$ TTs)
current (1,0,0)	1×1	none	1×1
2,1,1	2×2	5×5	4×4
2,0,1	2×2	3×3	4×4
3,1,1	3×3	5×5	5×5
3,-1,1	3×3	3×3	5×5

4.6.1.1 Energy resolution and turn-on curves

The choice of the size of the areas which determine the “trigger jets” has first been studied by looking at the energy resolution achieved, on samples of simulated events, with the following algorithms:

a) The ***R*** size is 0.6 x 0.6 (Figure 40a) and the trigger E_T is the E_T contained in the RoI – algorithm 3,0,0.
b) The ***R*** size is 0.4 x 0.4 (Figure 40b) and the trigger E_T is the E_T contained in the 0.8 x 0.8 region around the RoI – algorithm 2,1,1.
c) The ***R*** size is 1.0 x 1.0 (5x5 TTs) and the trigger E_T is the E_T contained in the RoI – algorithm 5,-1,0.

In each case, the algorithm illustrated in Figure 41 is used to find the local maxima ***R***. For each algorithm, the transverse energy seen by the trigger for 40 GeV jets is shown in Figure 43. This is to be compared with Figure 26, which shows the E_T seen by the current trigger. Clearly, any of the "sliding window" algorithms considerably improve the resolution of the trigger E_T. For the case of the 40 GeV jets studied here, the resolution improves from an rms of about 50% of the mean (for a fixed 0.2x0.2 η x ϕ trigger tower) to an rms of 30% of the mean (for a sliding window algorithm), and the average energy measured in the trigger tower increases from ~26% to 56-63% (depending on the specific algorithm).

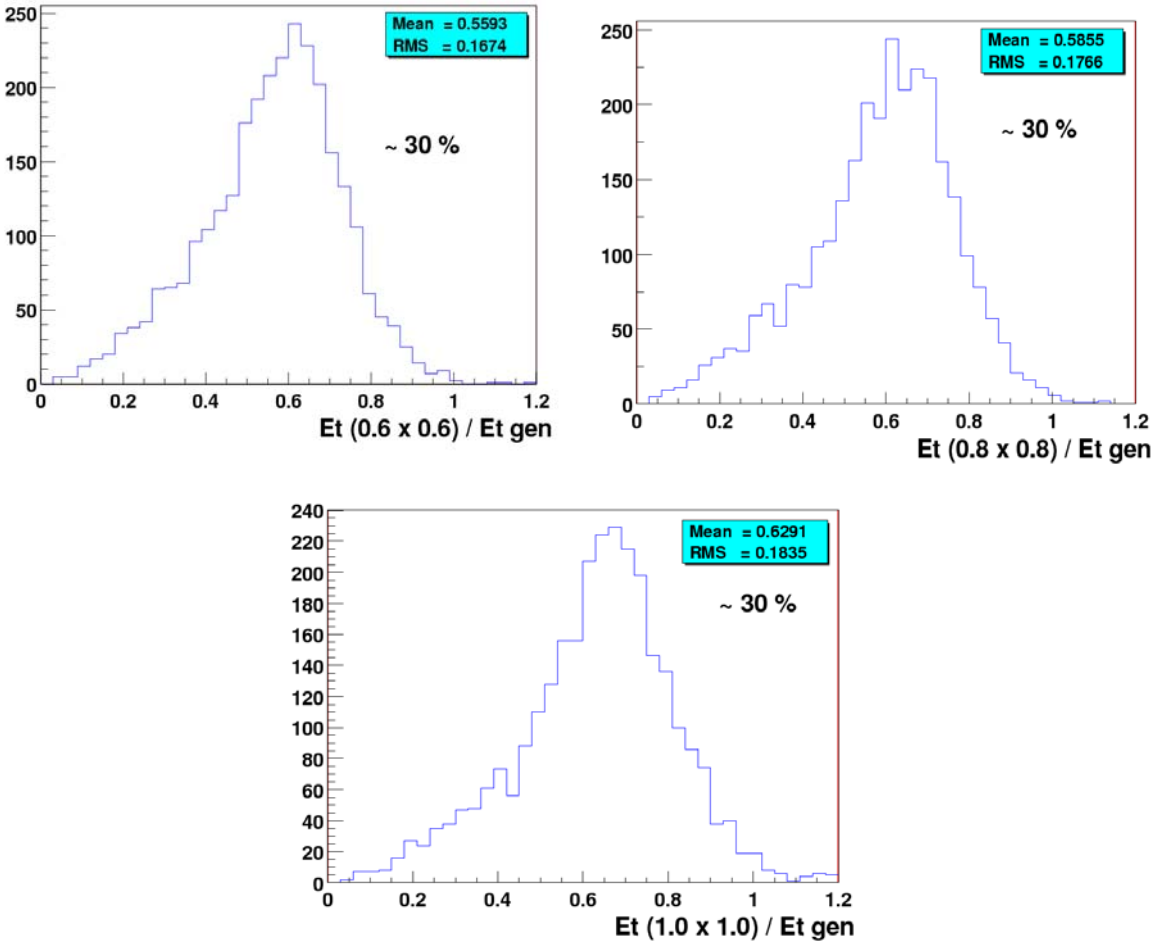

Figure 43. Ratio of the trigger E_T to the transverse energy of the generated jet, using three different algorithms to define the trigger jets. Only jets with $E_T \approx 40$ GeV are used here. The ratio of the rms to the mean of the distribution, the value 30%, is written on each plot.

Since the observed resolution is similar for all three algorithms considered, then the choice of the ***R*** definition (*i.e.* of the algorithm) is driven by other considerations including hardware implementation or additional performance studies.

The simulated trigger efficiency for the 2,1,1 (b) algorithm, with a threshold set at 10 GeV, is shown as a function of the generated E_T in Figure 44. The turn-on of the efficiency curve as a function of E_T is significantly faster than that of the current trigger, also shown in Figure 44 for two values of the threshold. With a 10 GeV threshold, an efficiency of 80% is obtained for jets with E_T larger than 25 GeV.

In order to understand which part of these new algorithms are providing the improvement (the sliding window or the increased trigger tower size), we have studied the gain in efficiency which is specifically due to the sliding window procedure by considering an algorithm where the TTs are clustered in fixed 4 x 4 towers (i.e. 0.8x0.8 in $\eta x \phi$), without any overlap in η or ϕ. The comparison of the "fixed" and "sliding" algorithms is shown in Figure 45. One observes a marked improvement for the "sliding" windows compared to the "fixed" towers, indicating that the added complexity of implementing sliding windows is warranted.

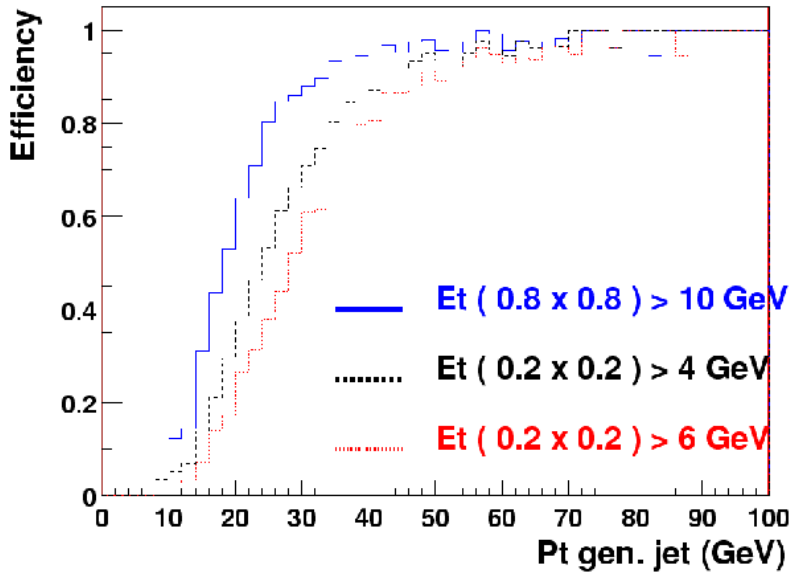

Figure 44. Trigger efficiency as a function of the transverse energy of the generated jet, for the (b) algorithm for E_T >10 GeV (the solid line) and for the current trigger (fixed trigger towers with thresholds of 4 and 6 GeV shown as dashed and dotted lines respectively).

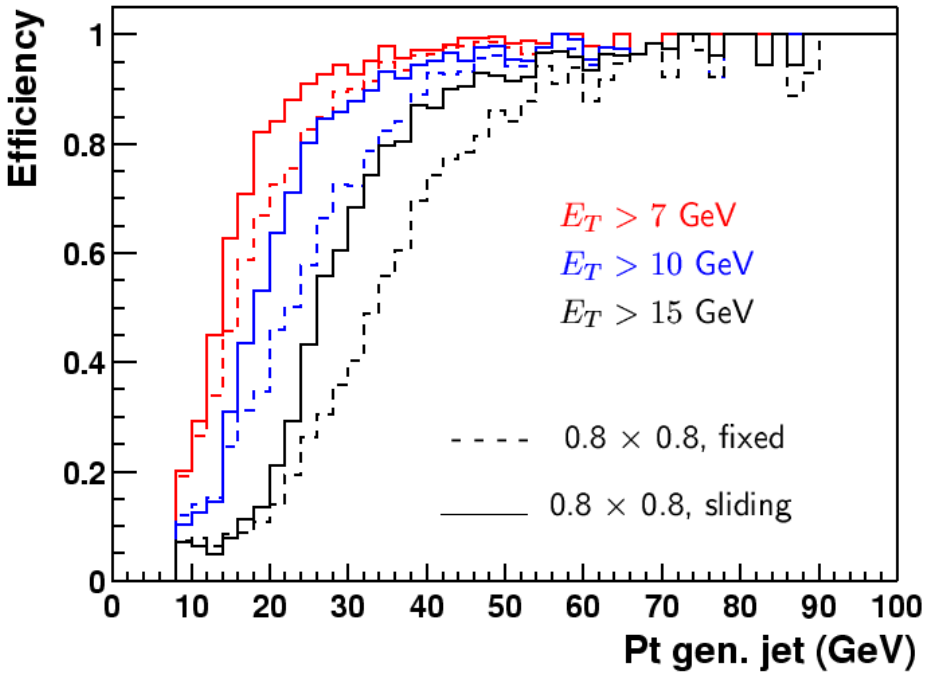

Figure 45. Trigger efficiencies as a function of the generated jet p_T for trigger thresholds E_T > 7GeV, 10 GeV and 15 GeV (curves from right to left respectively). The solid curves are for the 0.8 x 0.8 "sliding window" algorithm, and the dashed curves are for a fixed 0.8 x 0.8 trigger tower in $\eta x \phi$.

4.6.1.2 Position Resolution

Accurately reporting the position of trigger jets is an important task of the L1 calorimeter trigger. This information is used for matching to tracks in the L1 Cal-Trk match system. Poor resolution in the calorimeter position measurement then translates directly into worse background rejection in the Cal-Trk match. One of the main advantages of algorithms with 2×2 sized windows over those with windows of size 3×3 is that the 2×2 algorithms give less spurious shifts in cluster position and therefore better position resolution. This problem is illustrated for the 3,-1,1 algorithm in Figure 46, where energy deposited in a single tower in the presence of a small amount of noise causes two local maxima to be found (which have nearly identical cluster energies), both of which are offset from the position of the original energy deposition. This problem is avoided in the 2,1,1 algorithm.

The reason for the better behavior of 2×2 algorithms compared to 3×3 algorithms is that 2×2 windows reflect the asymmetry inherent in the declustering scheme (see Figure 41) if the "anchor" TT in the window is taken to be in the lower left corner.

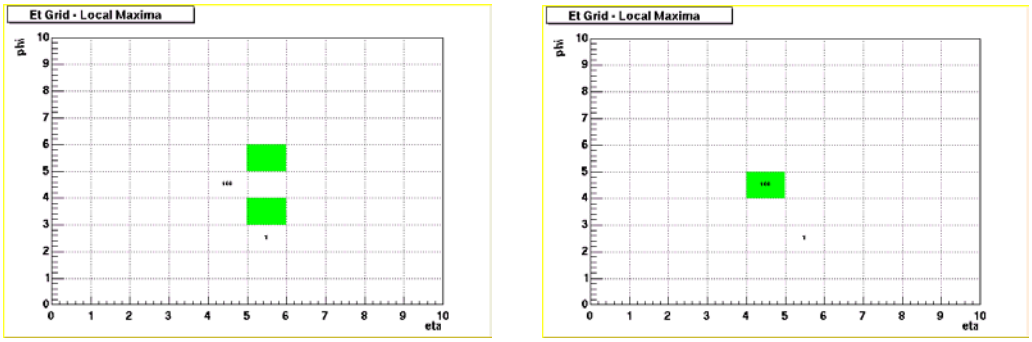

Figure 46: Results of various declustering algorithm for energy deposited in a single TT plus a small noise deposition in a neighboring TT. Positions of local maxima are shown shaded. The left hand plot uses the 3,-1,1 algorithm for declustering while the right hand plot uses the 2,1,1 algorithm.

4.6.1.3 Trigger jet multiplicities

The number of jets above a given threshold in E_T will be an important ingredient of any trigger menu. As is evident from Figure 46 some algorithms give spurious jets in the presence of noise in the calorimeter, particularly for the case of narrow energy deposition. Note that high energy electrons, which deposit their energy in one TT, will be reconstructed as trigger jets as well as EM clusters with the jet algorithm proposed. The 3,-1,1 algorithm, which has a declustering region of 3×3 windows, gives the largest probability to reconstruct spurious jets. All other algorithms attempted, including the 2,0,1 version, which also declusters in a 3×3 region, yield approximately the same performance.

To give a more quantitative estimate of the size of this effect, we compare the jet multiplicities obtained on simulated events using two algorithms: 3,-1,1 (declustering region of 3×3 TTs) and 3,0,1 (declustering region of 5×5 TTs).

The mean number of jets with E_T above a given threshold is shown in Figure 47, for a sample of simulated QCD events (upper plot), and for pair-produced top quarks which decay fully hadronically (lower plot) leading to high E_T jets. Both trigger algorithms lead to comparable multiplicities, especially when high E_T trigger jets are considered. The multiplicity of jets found by an offline cone algorithm of radius 0.5 is also shown in Figure 47 as the thin line. It is larger than the trigger jet multiplicity, as expected since the trigger jet E_T is not 100% of the reconstructed E_T.

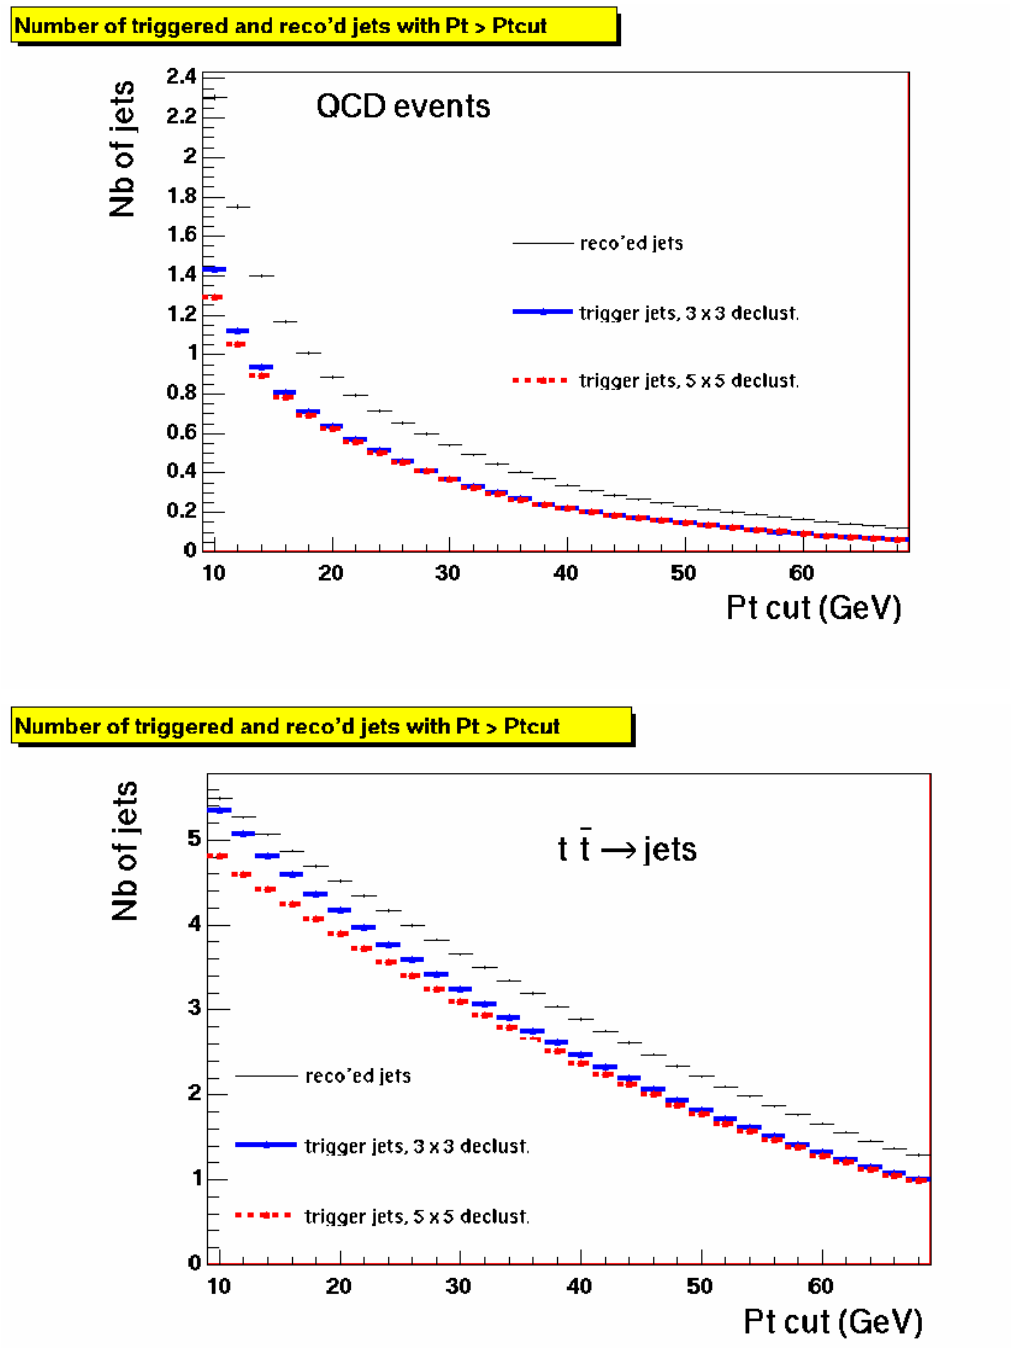

Figure 47: Multiplicities of jets with E_T above a given cut, as found by two trigger algorithms differing in the declustering procedure. The multiplicity of reconstructed jets is also shown.

From this study it is evident that a higher jet multiplicity is found, especially at low P_T for the 3×3 declustering region (algorithm 3,-1,1).

4.6.1.4 Rates and rejection improvements

In this section, we compare the performance of the sliding window and the existing trigger algorithms. We compare both of these algorithms' trigger efficiencies and the associated rates from QCD jet events as a function of trigger E_T.

Rates versus trigger efficiency on hard QCD events

In these studies we require that for the 2,1,1 sliding window algorithm there be at least one region of interest with a trigger E_T above threshold which varies from 5 to 40 GeV in steps of 1 GeV. Similarly, for the current trigger algorithm, we require at least one TT above threshold which varies from 2 GeV to 20 GeV in steps of 1 GeV. For both algorithms and for each threshold, we calculate the corresponding inclusive trigger rate and the efficiency to trigger on relatively hard QCD events, *i.e.* with parton p_T > 20GeV and p_T > 40GeV respectively. To simulate high luminosity running, we overlay additional minimum bias events (a mean of 2.5 or 5 additional minimum bias events) in the Monte Carlo sample used to calculate the rates and efficiencies. While the absolute rates may not be completely reliable given the approximate nature of the simulation, we believe that the relative rates are reliable estimators of the performance of the trigger algorithms. Focusing on the region of moderate rates and reasonable efficiencies, the results are plotted in Figure 48 where lower curves (open squares) in the plots are for the current trigger algorithm and the upper curve (solid circles) corresponds to the 2,1,1 sliding window algorithm. It is apparent from Figure 48 the sliding window algorithm can reduce the inclusive rate by a factor of 2 to 4 for any given efficiency. It is even more effective at higher luminosities (*i.e.* for the plots with 5 overlaid minimum bias events).

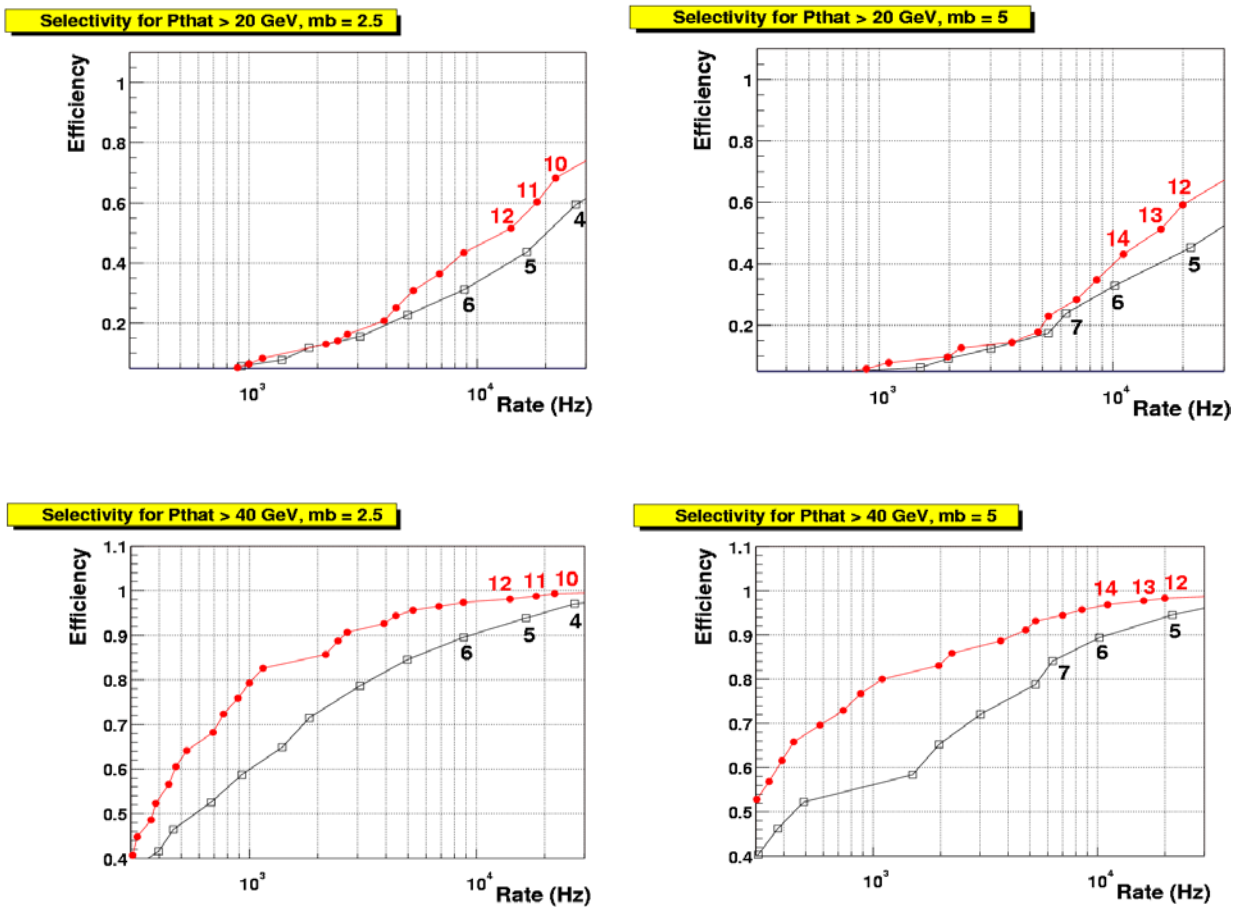

Figure 48. Trigger efficiency for events with parton p_T > 20 GeV (upper plots) and parton p_T > 40 GeV (lower plots) as a function of the inclusive trigger rate, for the (b) algorithm (solid circles) and the current algorithm (open squares). Each dot (solid circle or open square) on the curves corresponds to a different trigger threshold; the first few are labeled in GeV, and they continue in 1 GeV steps. The luminosity is $2x10^{32}$ cm^{-2} s^{-1} and the number of overlaid minimum bias (mb) events follows a Poisson distribution of mean equal to 2.5 (left hand plots) or to 5 (right hand plots).

Rates versus trigger efficiency on events with a large hadronic activity

In this section we study the performances of sliding algorithms on events which have a large number of jets in the final state. As an example we consider the case of pair produced top quarks which both decay fully hadronically. Other topologies with large jet multiplicities could arise from the production of squarks and/or gluinos.

Three sliding algorithms have been considered here:

i. The size of the regions of interest is 0.6x0.6 (i.e. 3x3 TTs); the trigger E_T is that of ***R***; the declustering is performed in a 5x5 window. This algorithm is labeled 3_0_0.

ii. As (i) but the trigger E_T is obtained by summing the E_T of ***R*** and the E_T of the closest neighboring TTs. This algorithm is labeled 3_0_1.

iii. As (ii) but the declustering is performed in a 3x3 window. This algorithm is labeled 3_m1_1.

In each case, the trigger condition requires that there be at least three trigger jets with E_T above a varying threshold. In addition, the E_T of the highest E_T jet should be above 40 GeV. A similar trigger condition has also been applied using the 0.2x0.2 TTs instead of the trigger jets; in this latter case the highest E_T TT should have E_T > 15 GeV. The inclusive QCD rate has been obtained as before, using QCD Monte Carlo events where a mean number of 7.5 minimum bias events has been overlaid. Figure 49 shows the resulting efficiencies and rates. Inclusive rates are shown here for a luminosity of 5×10^{32} cm^2 s^{-1} .

It can be seen that the three sliding algorithms considered lead to very similar performances. In particular, no noticeable difference is seen between algorithms 3_0_1 and 3_m1_1 (which differ by the declustering procedure only), as was seen in Section 4.6.1.2. The figure also shows that the performances of sliding algorithms are better than those of the current trigger system, also for events with many jets in the final state.

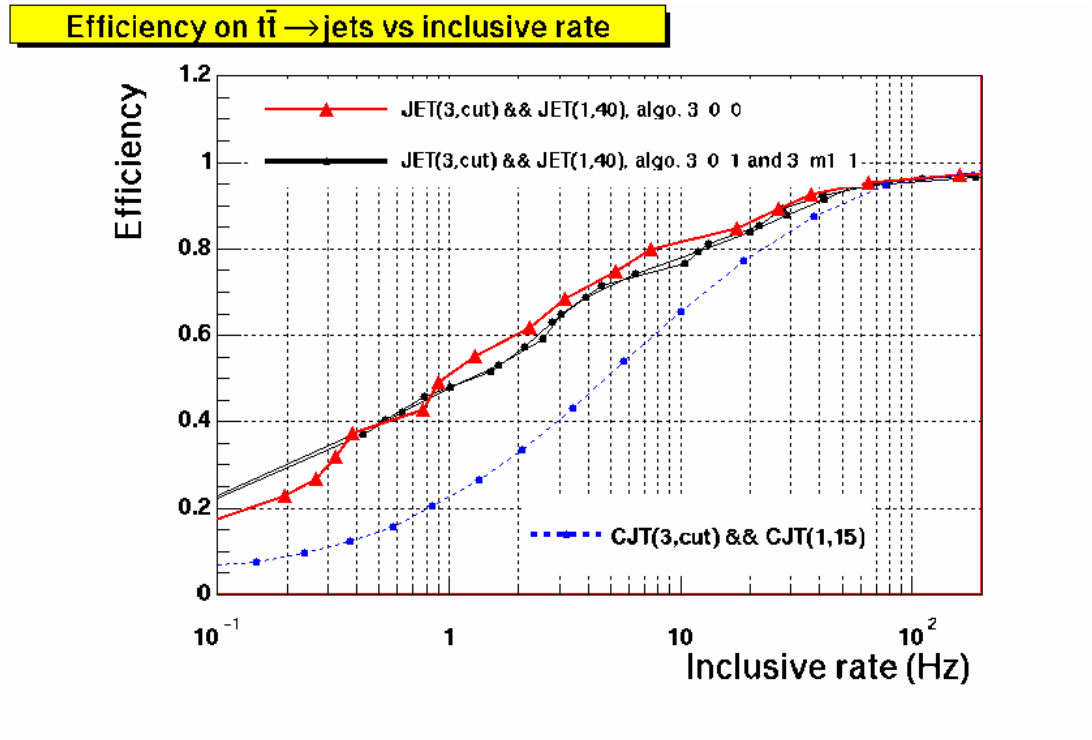

Figure 49. Trigger efficiency for simulated pair produced top quarks which both decay hadronically, as a function of the inclusive trigger rate, for various sliding window algorithms (full curves, solid circles and triangles), and using the current trigger towers (dashed curve, solid circles). The trigger condition for the sliding (current) algorithms requires at least three jets (TTs) with E_T above a varying threshold; the highest E_T jet (TT) must moreover satisfy E_T > 40 GeV (E_T > 15 GeV).

Rates versus trigger efficiency on "difficult" topologies

The improvement in jet triggering provided by the proposed algorithm is important for those physics processes that do not contain a high p_T lepton which in and of itself offers considerable rejection. Since the sliding window algorithm would be implemented in FPGA-type logic devices, it opens up the possibility of including further refinements in the level of trigger sophistication, well beyond simple counting of the number of towers above threshold. We have studied the trigger for two processes which demonstrate the gains to be expected from a sliding window trigger over the current trigger:

- The production of a Higgs boson in association with a $b\bar{b}$ pair. This process can have a significant cross-section in supersymmetric models with large tanβ, where the Yukawa coupling of the b quark is enhanced. Thus when the Higgs decays into two b quarks this leads to a 4b signature. The final state contains two hard jets (from the Higgs decay) accompanied by two much softer jets. Such events could easily be separated from the QCD background in off-line analyses using b-tagging. But it will be challenging to efficiently trigger on these events while retaining low inclusive trigger rates.

- The associated production of a Higgs with a Z boson, followed by $H \to b\bar{b}$ and $Z \to \nu\bar{\nu}$. With the current algorithm, these events could be triggered on using a di-jet + missing energy requirement. The threshold on the missing energy could be lowered if a more selective jet trigger were available.

Figure 50 shows the efficiency versus inclusive rate for these two processes, where three different trigger conditions are used:

1. At least two fixed trigger towers of 0.2 x 0.2 above a given threshold (dotted curves, open squares).
2. At least one TT above 10 GeV and two TT above a given threshold (dot-dash curve, solid stars).
3. At least two "trigger jets" whose summed trigger E_T's are above a given threshold (solid curve, solid circles).

The algorithm b) has been used here. It can be seen that the third condition is the most efficient for selecting signal with high efficiency but low rates from QCD jet processes.

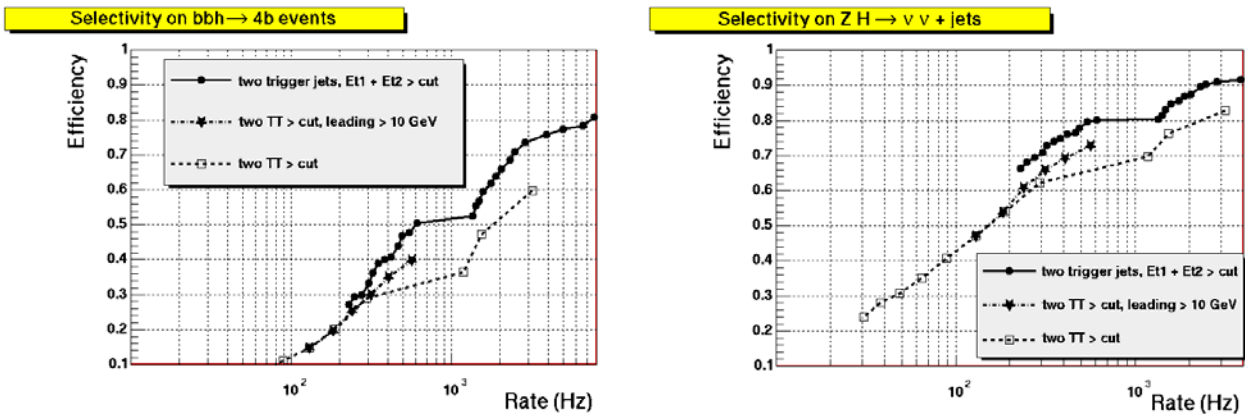

Figure 50. Efficiency to trigger on bbh (left) and ZH (right) events as a function of the inclusive rate. The three conditions shown require: at least two TT above a threshold (dotted, open squares), at least one TT above 10 GeV and two TT above a threshold (dot-dash, solid stars), at least two trigger jets such that the sum of their trigger E_T's is above a given threshold (solid circles).

4.6.1.5 Including ICR in trigger jets algorithms

In the current trigger system, jet energy resolutions in the inter-cryostat region (ICR) are degraded because energy from the ICR detectors is not included in the TTs. We have studied the effect of adding this energy in to TTs in the ICR. See Section 4.6.4 for more details about ICR energy. Three possibilities were considered:

1. TTs do not include energy from ICR detectors.
2. TTs do not include energy from ICR detectors, but thresholds for jets in the ICR were adjusted to take this into account.
3. ICR detector energies were added to their respective TTs.

Including ICR detector energy in TTs in the ICR gives modest improvements over the cases where this energy was not included and where adjusted thresholds were used in the region.

4.6.2 Electron algorithms

The EM cluster (electron) algorithm we have chosen to implement as a baseline is shown in Figure 51.

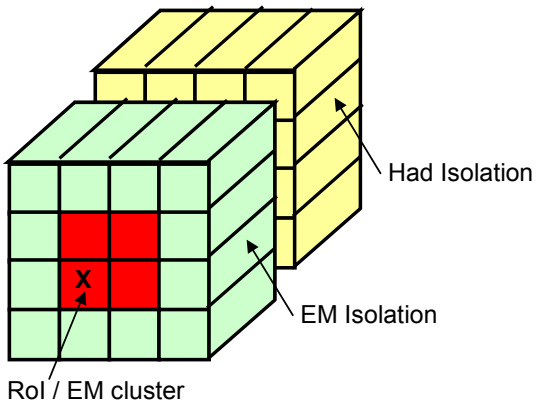

Figure 51: A schematic diagram of the baseline electron algorithm.

Its main features are:

1. A 2×2 window size (using only EM E_T's) with declustering performed over a region of 5×5 windows and trigger cluster energy defined over the same region as the window – the 2,1,0 algorithm.
2. Electromagnetic isolation enforced by requiring the energy in the ring of TTs surrounding the window to be less than a fixed amount.
3. Hadronic isolation performed by requiring that the energy in the 4×4 region directly behind the window be less than a specified fraction of the cluster energy. Note: in order to fit well into the electronics implementation of this algorithm, only fractions that correspond to powers of 2 are allowed in the comparison. In the following 1/8 is the generally used value.

4.6.2.1 Improvements

Studies to optimize electron algorithm parameters are still ongoing, although preliminary results indicate that the version described above yields acceptable efficiencies and background rejections. This version is therefore being used in the baseline design of the system. Since it is the most complicated of the

electron algorithms that we are considering, this should lead to a conservative design.

4.6.3 Tau algorithms

With some refinements, the sliding window algorithms presented in Section 4.6.1 could lead to some sensitivity to the process $gg \rightarrow H \rightarrow \tau^+\tau^-$ at the trigger level. This could be achieved by exploiting the fact that τ jets are narrower than "standard" jets. In its most basic form, such an algorithm would be simple to implement using information already calculated for the jet and EM clusters. We are studying an algorithm with the following characteristics:

1. A 2×2 window size (using EM+H ET's) with declustering performed over a region of 5×5 windows and trigger cluster energy defined over the same region as the window - the 2,1,0 algorithm. Window formation and declustering would be performed as part of the jet algorithm.
2. Narrow jets (isolation) defined by cutting on the ratio of the EM+H energy in the 2×2 jet cluster to that in the 4×4 region centered on the cluster. Both of these sums are available as part of the jet algorithm calculations.

4.6.3.1 Transverse isolation of τ jets

We consider here the 2,1,1 sliding window algorithm described in Section 4.6.1, where the size of regions of interest is 0.4x0.4 (2x2 TTs), while the size of trigger jets is 0.8x0.8 (4x4 TTs). We compare the ratio of the ***R*** E_T to the trigger E_T, for τ jets coming from $gg \rightarrow H \rightarrow \tau^+\tau^-$ and for jets coming from QCD processes. As shown in Figure 52, QCD jets become more and more collimated as their E_T increases, but the ratio of the "core E_T" to the trigger jet E_T (called the "core fraction") remains a powerful variable to discriminate between τ jets and QCD jets.

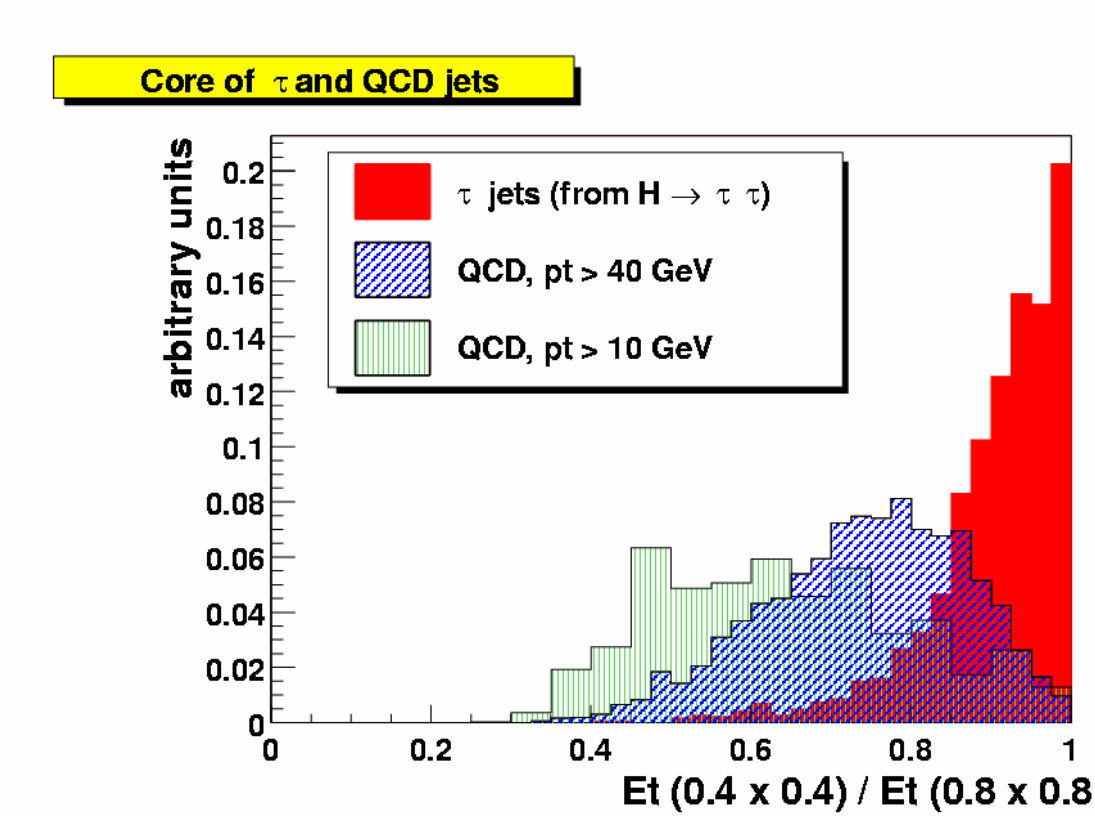

Figure 52: Ratio of the R E_T to the trigger E_T, for the sliding window algorithm (b). The ratio is shown for τ jets coming from a Higgs decay (full histogram), and for jets coming from QCD processes (hashed histograms).

4.6.3.2 Rates and rejection improvement

This can be exploited by defining a specific trigger condition, which requires at least two jets whose summed trigger E_T's is above a threshold, and for which the core faction is above 85%. As can be seen in Figure 53, it seems possible to have a reasonable efficiency on the signal (70 %) while maintaining the inclusive rate below 300 Hz. The figure also shows that such an algorithm reduces the inclusive rate by a factor of about 3, compared to the current trigger system.

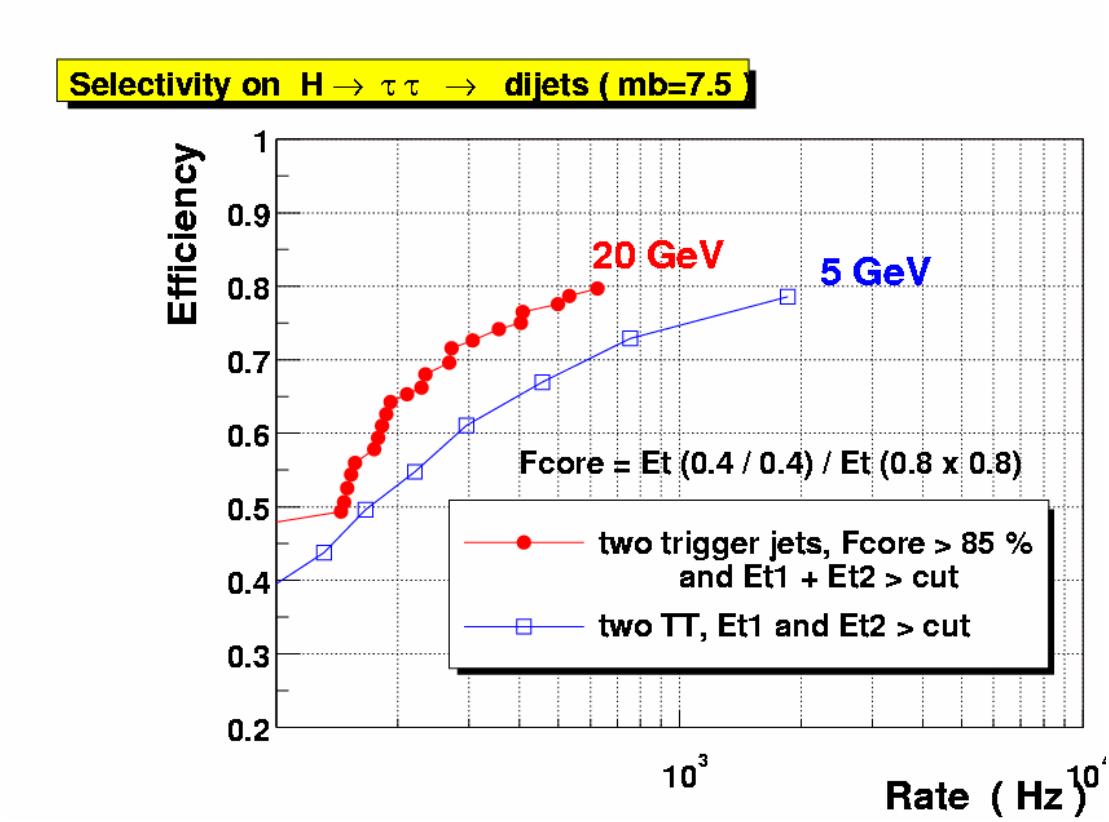

Figure 53: Efficiency to trigger on gg →H →ττ events as a function of the inclusive QCD rate, for: (closed circles) the sliding window algorithm (b), when requiring at least two jets whose summed E_T is above a varying threshold, and whose core fraction is above 85%; (open squares) the current trigger system, requiring two TTs whose summed E_T is above a varying threshold. The inclusive rates shown here correspond to a luminosity of 2 x 10^{32} cm^2 s^{-1}.

4.6.4 Global sums

The region around 0.8<|η|<1.5, known as the inter cryostat region (ICR) encompasses the transition from showers contained within the CC and showers contained within the EC. There is a gap in EM coverage and a major thinning of FH coverage in this area. Since these are the layers which comprise standard trigger towers, there is a major degradation in Level 1 calorimeter response and resolution in this region. This is exacerbated by the presence of significant dead material in the solenoid in Run2. To aid in recovering the energy losses in the ICR region, we use the intercryostat detectors (ICD), which consists of scintillators located in the gap between the calorimeter cryostats, and the "massless gaps" (MG) which consist of the front sections of the endcap calorimeter that have no absorber in front. In this section we study ways to improve the energy measurement at the trigger level.

4.6.4.1 Concept & performance

Global tower E_T sums such as missing E_T or scalar E_T, while very useful, suffer from several significant problems at the L1 trigger. There are two significant issues: first is that the ICR sampling layers are not available in the

calculation at Level 1; second is that the imprecision of the tower E_T's gets compounded for global sums, resulting in significantly degraded effectiveness. This is particularly true in a multiple interaction environment. There are two possible solutions to these problems. First we can take advantage of work done for Run IIa to make the ICR layers available at Level 2 and add these towers back into the global sums at Level 1 in Run IIb. Second, we can develop a scheme which discriminates towers which are from multiple interactions and avoids adding them into the sum.

Simulations of single pions and jets in this region indicate that the energy scale in this region goes as low as 40% of the CC/EC scale (as shown in Figure 54), and the resolution is as bad as 6 times worse than in the CC or EC (as shown in Figure 55). These results are very consistent with findings from Run1 Level 1 missing E_T analyses (see Figure 56). One of the major results of this deficiency is that the efficiency and rejection of a Level 1 missing E_T selection are noticeably degraded. These simulations also indicate that adding ICD and MG cells into trigger towers can improve the scale by a factor of 2, while optimizing the resolution by a factor of 3.

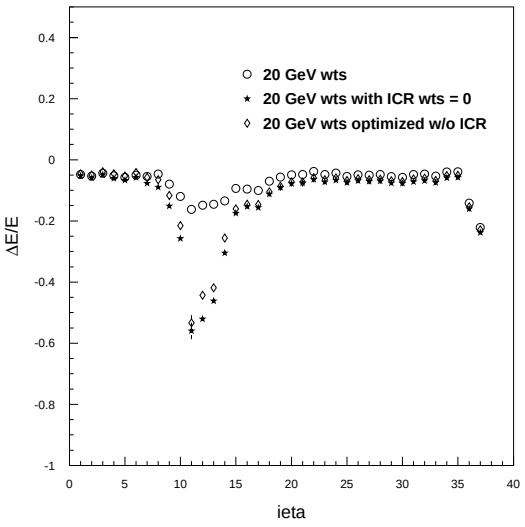

Figure 54. The relative calorimeter energy response in the ICR region for incident 20 GeV pions as a function of □x10. The stars are the response if the ICR weights are set to zero, the open diamonds are the response if the ICR energies are ignored and the remaining calorimeter weights are re-optimized, and the open circles are the response when the ICR region is included and the weights are optimized.

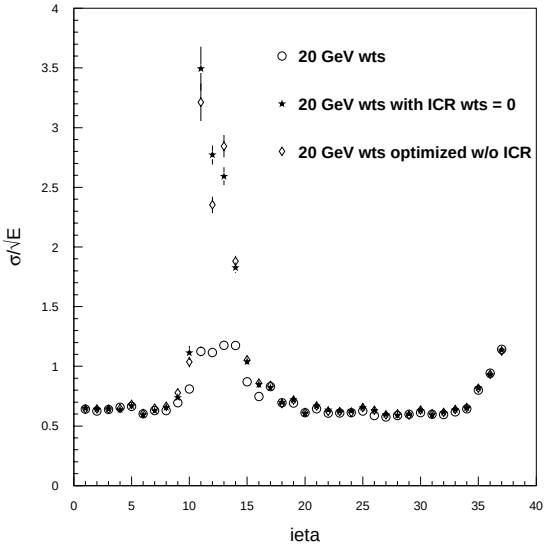

Figure 55. The calorimeter energy resolution in the ICR region for incident 20 GeV pions as a function of □x10. The stars are the response if the ICR weights are set to zero, the open diamonds are the response if the ICR energies are ignored and the remaining calorimeter weights are re-optimized, and the open circles are the response when the ICR region is included and the weights are optimized.

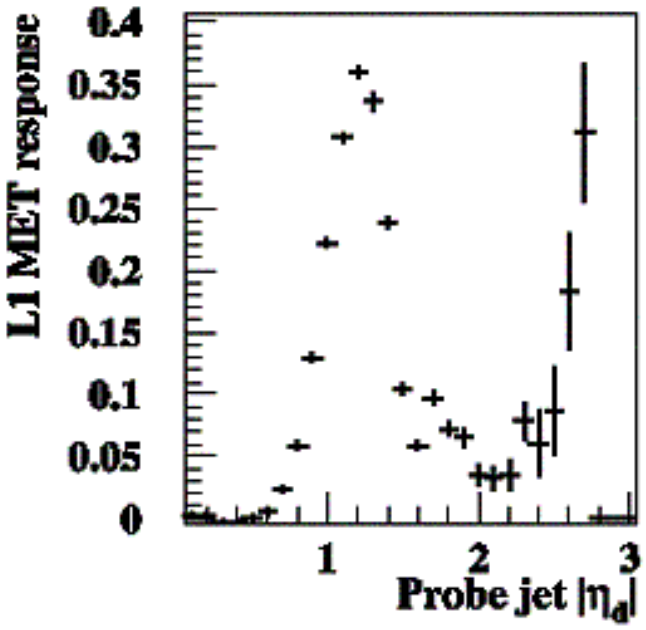

Figure 56. The L1 missing E_T response as a function of η for 85 GeV jets using the Run 1 DØ detector simulation.

4.6.4.2 Simulation results

In principle, it is straightforward to estimate the effect of the ICD and MG to the missing E_T calculation. However our present simulations do not yet fully address a number of issues (including a proper treatment of trigger tower sampling weights and the verification of the modeling and calibration for the ICR). To estimate the effect of adding the ICR detectors into the missing E_T, we therefore consider the fact that in the region of $1.0<|\eta|<1.4$, the sampling weight simulations indicate approximately half of the energy will be deposited in the EM+FH, and the other half in the ICD+MG. As a crude estimate of the magnitude of the effect of adding the ICR layers, we will merely consider the missing E_T measurement with and without the EM+FH layers in this region and assume the ICR improvement will be similar. Although the sample used for this calculation is a QCD sample with jet p_T > 20 GeV and 0 minimum bias events overlaid, for historical reasons it is a different sample than that mentioned in the rest of this document with the same specifications. The missing E_T mean and rms in this sample behave as follows:

if remove all ICR TTs: μ/rms = 6.7 GeV / 4.8 GeV

if only use EM+FH TTs: μ/rms = 5.5 GeV / 3.9 GeV

The number of events passing various Level 1 missing E_T cuts in this sample are shown in Table 7.

Table 7. Events passing L1 missing E_T cuts when the ICR energy is included and when it is removed from the trigger towers.

L1 ME_T	Without ICR	With ICR
> 5GeV	948	766
> 10 GeV	337	185
>15 GeV	95	40
> 20 GeV	37	11
> 25 GeV	9	4

Thus, the region is important to the missing E_T calculation and the rates of passing 15 or 20 GeV selection can change by factors of around 2.5-3. A proper treatment of the gains from adding in the ICD and MG, however, will have to await a satisfactory treatment of the relative weights of various layers.

4.6.4.3 Improving Missing E_T for Multiple interaction Events

Our experience in Run1 indicated the Level 1 missing E_T to be very sensitive to the number of multiple interactions. This results from several factors, including the fact that the fundamental trigger tower fractional energy resolution is poor, especially for very low E_T towers, and the numbers of these towers increases substantially with the number of multiple interactions. As a result, we have explored three ways in which we might improve the missing E_T resolution to reduce this problem in Run IIb.

First, we varied the low threshold on the E_T of towers going into the global sum. In Run1, this threshold was 0.5 GeV and was not studied in detail in the light of multiple interactions. Again, we have used the QCD p_T>2 GeV and p_T>20 GeV samples with 0 minimum bias (mb) events overlaid, a 5mb overlay, and a 10mb overlay. We have used the $t\bar{t}$ sample with 2.5 mb overlays for the signal. If we calculate the missing E_T mean and the rms in these samples for various E_T thresholds, we find the results shown in Table 8.

Table 8. Change in the means and rms for the missing E_T for background (QCD) and signal ($t\bar{t}$) samples as a function of the trigger tower (TT) threshold. A selection of 1.5 GeV on trigger towers removes most of the multiple interaction variation for the QCD samples, while having little effect on the signal top sample.

ME_T calc	2 GeV QCD (μ/rms) in GeV	20 GeV QCD 0mb (μ/rms) in GeV	2 GeV QCD 10 mb (μ/rms) in GeV	20 GeV QCD 5mb (μ/rms) in GeV	ttbar (μ/rms) in GeV
TT>0.5GeV	1.0/1.0	5.1/3.8	3.1/2.2	6.5/4.2	35.9/25.4
TT>1GeV	0.6/0.9	5.2/3.9	2.3/1.9	5.8/4.0	35.4/24.7
TT>1.5GeV	0.3/0.7	5.3/4.1	1.6/1.9	5.6/4.0	35.0/24.1
TT>2GeV	0.1/0.6	5.2/4.2	1.0/1.7	5.4/4.2	34.6/23.6

The error on the mean and RMS for the QCD samples is approximately 0.1 GeV. The cut of 2GeV reduces the mean of the QCD sample noticeably. If we consider the 20 GeV sample, the trigger tower cut of 1.5 GeV provides a 20% to 30% lower pass rate for moderate missing E_T selections. Although scalar E_T is generally considered a poor variable at Level 1 because of its sensitivity to multiple interactions, we have studied its mean and rms (see Table 9) for the same thresholds to see what is happening:

Table 9. Change in the means and rms for the E_T scalar sum for background (QCD) and signal (ttbar) samples as a function of the trigger tower (TT) threshold.

Sum E_T calc	2 GeV QCD 0.7 mb(μ/rms) in GeV	QCD 0 mb (μ/rms) in GeV	2GeV QCD 0.7 mb(μ/rms) in GeV	QCD 5mb (μ/rms) in GeV	ttbar (μ/rms) in GeV
TT>0.5GeV	2.9/3.3	23.5/13.0	21.2/18.1	57.7/39.3	179.7/68.8
TT>1GeV	0.8/1.5	17.9/11.9	6.5/7.1	26.6/15.8	161.1/66.4
TT>1.5GeV	0.3/1.1	14.7/11.4	2.8/4.2	18.0/12.5	151/64.9
TT>2GeV	0.2/0.8	12.5/11.1	1.5/3.1	14.2/11.6	143.6/63.8

Comparison of the two QCD samples indicates that low thresholds let in an enormous amount of energy which has nothing to do with the hard scatter interaction.

Because the typical low p_T QCD event E_T is distributed flat in η, we might not expect a degradation in the global sum behavior from including forward trigger towers in the calculation of these quantities. In fact, when looking in simulated events even with large numbers of multiple interactions, one finds very little transverse energy in this region. However, our experience in Run1 indicated strongly that use of forward towers (i.e. those around $|\eta|$ ~3 or more) substantially degraded the missing E_T behavior. This was especially true in a multiple interaction environment. As a result, we suspect strongly that there is a benefit from being able to easily select what the range is for the calculation, or

perhaps include the η parameter into a weighting scheme with the trigger tower E_T. This requires further study only possible once data is available.

Another concern for the missing E_T measurement involves the truncation of trigger tower E_T's into 0.5 GeV bins. Since one to two hundred towers are typically added into the Missing E_T, this resolution loss can start to be noticeable. Taking the QCD p_T>20 GeV sample with minimum bias overlay of 0 and 1648 events, we can use the simulator described above in the ICR discussion and toggle truncation on and off. The results are shown in Table 10.

Table 10. Comparison of the effect of TT truncation on the missing E_T. The table lists the number of events (out of a sample of 1648, QCD with p_T> 20GeV and no minimum bias overlaid events) that pass the listed missing E_T thresholds.

Missing E_T	no truncation	no truncation, TT>0.5GeV	with truncation
>5 GeV	947	868	766
>10 GeV	309	261	185
>15 GeV	76	51	40
>20 GeV	22	17	11
>25 GeV	7	5	4

The first column indicates truncation turned off and no threshold applied to trigger towers. The second column also has no truncation and zeros out all towers with E_T <0.5. The third column employs the normal 0.5 GeV truncation. Since truncation lowers tower E_T's only to the next lowest 0.5 GeV increment, it effectively deweights all of the poorly measured E_T in low E_T towers. In fact, if we consider the QCD Pt>20 GeV sample with 5mb already discussed, the missing E_T mean and rms are mildly improved over the straight 1.5GeV threshold by a simple weighting scheme. If we choose weights of 5%, 25%, and 75% for E_T = 0.5, 1.0, and 1.5 GeV, respectively, we find the results shown in Table 11.

Table 11. Comparison of simple TT threshold vs. weighting scheme for 20GeV QCD jet sample.

Scheme	μ (GeV)	rms
TT E_T>1.5 GeV:	5.41	4.20
Weighted TT:	5.41	3.96

Because the trigger tower threshold seems to be the simplest solution that shows progress, and the weighting also seems to help, one might ask whether rejecting low E_T towers unless they are near significant neighbors might help. Looking again in the 5mb QCD sample at missing E_T means and sigmas, we find the results shown in Table 12. These results point to a significant degradation in missing the E_T mean and resolution.

Table 12. Comparison of effect of rejection low ET towers unless they are near trigger towers (NN) with significant energy deposits.

Cut	μ (GeV)	rms (GeV)
None :	6.45	4.17
NN E_T > 0.5 GeV:	6.45	4.37
NN E_T > 1.5 GeV:	6.56	4.37
NN E_T > 3.0 GeV:	6.72	4.86
NN E_T > 10 GeV:	5.62	4.57
NN E_T > 1k GeV:	5.41	4.20

4.6.4.4 Conclusions

In this section, we have explored several different ways to improve the calorimeter missing E_T measurement at Level 1. Studies leading to the optimization of the Run IIa trigger have indicated a large improvement in the scale and resolution of jets in this region if the ICD and MG are used. Although our current simulation samples do not have a proper treatment of this region, a crude estimate indicates that this amount of energy should have a noticeable improvement on the missing E_T resolution.

Several attempts were also made to improve the behavior of missing E_T in a multiple interaction environment. The most promising appears to be a simple tightening of the E_T threshold on a trigger tower to around 1.5 GeV which would reduce the background by around 20% in our QCD sample. We take this to be our baseline design. The actual degradation in the real data may be larger than we see here, however, and the corresponding gain may also increase. We will be in a better position to evaluate this when we have reliable data at various luminosities. There is some evidence that a weighting scheme would provide further benefits.

4.7 L1 Calorimeter Trigger Implementation

4.7.1 Constraints

Because the L1 calorimeter system needs to be integrated into the existing DØ DAQ system it must obey several constraints.

4.7.1.1 Existing interfaces

The interfaces of the new system to the existing hardware should be compatible. In particular the new system must interface to the input pickoff signals, the L1 framework, the L2 and L3 data, the clock, and the timing and control systems.

The layout of the existing input signal cables places a special constraint on the new system. Moving these cables from their current locations in the Moving Counting House would require an enormous effort. Physical locations of the

boards in the new system will have to be adapted to the locations of the existing cables.

4.7.1.2 L1 Latency

The total L1 trigger latency is 4.2 μsec. After accounting for all transit times and front end processing, the maximum time remaining for complete signal processing, which includes digitization, filtering and the processing of the cluster algorithms is less than 2.7 μsec.

Additional, constraints are placed on the latency of the new system by the requirement of transmitting calorimeter clusters to the Cal-Track Matching system in time for it to meet the total L1 latency requirement. To match the arrival time of tracks from the L1 track trigger at the Cal-Track Match cards the calorimeter trigger must send out its clusters within 1500 ns of the beam crossing.

4.7.2 L1 Calorimeter Trigger Architectural Overview

A block diagram of the new L1 calorimeter trigger system is shown in Figure 57.

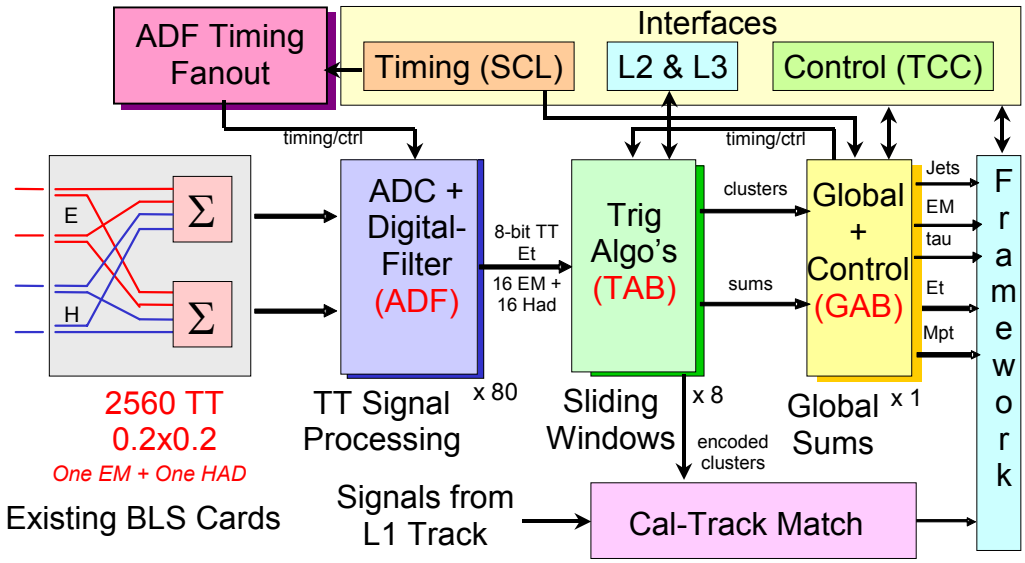

Figure 57. Block diagram of L1 calorimeter trigger, with the baseline subtractor (BLS) cards shown at left

The main elements of the system are listed below.

- ADC-Digital-Filter Boards (ADF) that receive analog TT signals from the BLS cards, digitize them, convert from energy to transverse energy (E_T) and perform the digital filtering to associate energy with the correct bunch

crossing. Each of these boards deals with signals from 16 EM TTs and 16 H TTs.

- ADF Timing Fanout boards that send timing signals coming from the trigger framework to the ADF cards.
- Trigger Algorithm Boards (TAB) that receive TT transverse energies from the ADF boards, produce EM and jet cluster E_T's using the sliding windows algorithm and begin the global summing process that will yield scalar summed transverse energy ($E_{T,total}$) and missing transverse energy (Mp_T). Outputs will also be provided at this level for data transmission to L2/L3 and to the Cal-Track Match system.
- A Global Algorithm Board (GAB) that receives data from the TABs and produces the final $E_{T,total}$ and Mp_T, as well as providing an interface to the DØ Trigger Framework and a timing fanout. One GAB is required for the system. It will be housed in the same crate as the TABs to facilitate communication between them.

These new electronics and associated communications cards will be housed in VME crates in two to three racks, replacing the present ADC and logic cards located in 10 racks.

Detailed design work has started on all elements of the L1 calorimeter trigger system. We describe below the various elements of the proposed calorimeter trigger.

4.7.3 Trigger Tower Mapping

The 2560 EM and Had Trigger Towers (TT) map to an $\eta\times\phi$ grid of 40x32 cells with 0.2x0.2 extent. This grid extends from –4.0 to 4.0 in η and from 0 to 2π in ϕ. Cells are referenced by η- and ϕ-indices (*ieta* and *iphi*) which run from –20-20 (excluding 0) and from 1-32, respectively. This grid has been modified, for Run 2, from a strict geographic mapping at the extreme values of *ieta* ($\pm$19 and $\pm$20) to allow inclusion of ICR detectors (the ICD and the MG) in the trigger.

The mapping of TTs in η is given in Table 13. All mapping in ϕ corresponds to the physical location of the TT in azimuthal angle. In the table, the column marked "Cable" signifies the labeling scheme for the cable carrying the data, while in the "Comments" column, the *EM* or *Had* refers to whether the relevant TTs appear physically in the EM or Hadronic parts of a shower and to which TTs the ICR energies should be added.

Table 13: Trigger Tower mapping scheme.

TT |*ieta*|	Cable	Detector η and ϕ	Comments
all	EM & Had	ϕ = (*iphi* – 0.5) x $2\pi/32$	
$\leq$ 16	EM & Had	η = (sign *ieta*) x (|*ieta*| - 0.5) x 0.2	*EM:* include only EM cells *Had:* include only FH cells
17	EM	$<\eta>$ = 3.45	*EM:* include only EM cells
	Had	$<\eta>$ = 3.45	*Had:* include only FH cells
18	EM	$<\eta>$ = 3.9	*EM:* include only EM cells
	Had	$<\eta>$ = 4.1	*Had:* include only FH cells
19	EM	$<\eta>$ = 0.95	MG $\rightarrow$ Cal *ieta* = 5 Had
	Had	$<\eta>$ = 1.25	ICD&MG $\rightarrow$ Cal *ieta* = 7 EM
20	EM	$<\eta>$ = 0.75	MG $\rightarrow$ Cal *ieta* = 4 Had
	Had	$<\eta>$ = 1.05	ICD&MG $\rightarrow$ Cal *ieta* = 6 Had

4.7.4 ADF System design and Implementation

The ADF system comprises the following components:

- ADF boards. These make the conversion of trigger pickoff analog signals to digital format and perform digital filtering to compute for each channel and for each beam crossing an 8 bit E_T calibrated energy value. These digital values are sent to the TAB system via high-speed links.
- ADF crates and controller interfaces. Each crate houses a number of ADF boards and includes an interface to the Trigger Control Computer (TCC). This slow path is used for downloading, calibration and monitoring.
- A timing distribution card. This card is connected to a Serial Control Link (SCL) receiver and is in charge of distributing the necessary clocks, synchronization and control signals within the ADF system.

In order to be able to test prototype ADF cards in spy-mode on the current DØ Run IIa setup, analog signal splitter cards are also designed. These duplicate the analog signals of several BLS's and connect to the CTFEs and the ADF cards being tested.

4.7.5 ADF Card Description

4.7.5.1 ADC precision, rate and analog section

Analog signals are digitized with 10-bit precision. The input voltage range is identical for all channels. This scheme is adequate to guarantee 0.25 GeV of resolution and a 62 GeV scale in E_T for all η values. The proposed conversion frequency is 30.28 MHz (i.e. BC x 4). This rate is a good trade-off between the

aim of a short conversion latency and the cost of a faster analog to digital converter. The ADC that was chosen is Analog Devices AD9218 (dual 40 Mbps 3V 10-bit ADC). Conversion latency is 165 ns when operated at 30.28 MHz.

The analog section of the ADF card includes:

- A differential receiver whose input impedance matches that of the driving cable and that provides an attenuation factor to bring trigger pickoff signals in the desired range;
- Circuitry for baseline subtraction to adjust the offset of the previous signal; a digital to analog converter produces a static correction voltage in a programmable way;
- The analog to digital converter previously mentioned;
- An anti-aliasing filter (2nd order Butterworth filter). The proposed cutoff frequency for this filter is 7.57 MHz given the fact that the spectrum of the signals to process does not have much content above that frequency. Tests will be conducted on the prototype and filter parameters will be modified if needed.

The device selected for the differential receiver is Analog Devices low distortion differential ADC driver AD8138. The device for the zero-offset digital to analog converter is Maxim octal 12-bit serial DAC MAX 5306. The anti-aliasing filter is implemented with passive RC components placed on the feedback loops of the differential receiver and at the input of the analog to digital converter.

4.7.5.2 Digital filters

The digital processing algorithm for trigger pickoff signals is a matched filter followed by a 3-point peak detector. Only static compensation of the baseline is performed. The matched filter operates at a rate of BC x 2. Decimation by a factor of 2 on the stream of values converted by the ADC is done prior to filtering. Selecting on a per tower basis the samples to process allows implementing some coarse compensation of delays. In addition, each channel can be individually programmed to have its analog input sampled on the rising edge or the falling edge of the common BC x 4 ADC sampling clock. Hence, the sampling phase for each channel can be adjusted in steps of 16.5 ns (i.e. 1:8 of the BC period). The filter comprises up to 8 taps. Implementing a lower number of taps is achieved by setting to 0 unused coefficients. Input samples are 10-bit unsigned numbers while coefficients are unsigned 6-bit numbers. Convolution is accumulated with the full 19-bit precision and 16-bits are kept for peak detection. The output of the peak detector is decimated by a factor of 2 (to produce only one output result per beam crossing) and results are truncated to 11-bit (1 overflow bit + 10 significant bits). These 11-bits are used to address a look-up table that implements the final E_T conversion (including clipping and saturation). The final result is an 8-bit quantity per channel and per beam crossing. All decimation parameters, coefficient values and look-up table content are programmable at run-time.

4.7.5.3 Raw samples and intermediate results buffering

Each channel keeps in circular memories the latest samples converted, the outputs of the filter (before peak detection) and the final E_T results. In the current design, the per-channel buffer uses a single 1Kx18bit block of dual-ported static RAM. Half of this memory (512 18-bit words) is used to hold the 18 most significant bits of the latest outputs of the digital filter. The other half of this memory is used to store the 512 latest raw ADC samples (10 bits) and final E_T results (8 bits). Raw ADC samples come at BC x 4. Because we need to time-share the input bus of the SRAM, it is clocked at BC x 8. A side effect is that each output of the filter (computed at BC X 2) is written twice, and each final output result (one per BC) is written four times. This is indeed not an issue because the depth of these history buffers translates to 128 BCs (i.e. 17 us), which is largely sufficient to cope with the latency of the Level 1 Trigger. The content of these buffers can be frozen under various circumstances (typically a Level 1 trigger) as explained in the next section.

A second use of the history buffers is to run the ADF card in test mode without analog signal inputs. For each channel, the history buffer can be pre-loaded with a sequence of raw samples values that are then played at the input of the digital filter. This mode of operation is useful for standalone functional tests, to check the correctness of computations with a known input pattern, etc.

4.7.5.4 Digital output to Trigger Algorithm Board

In order to simplify the design of the TAB system, it seems preferable to make the necessary duplication of data at the level of each ADF card, although this increases the number of cables between the two sub-systems. Each ADF card shall therefore include three identical output links. The net bandwidth per output link is 242.24 MByte/s, corresponding to 32 8-bit E_T values sent every 132 ns. In addition to that data, each ADF will also transmit the content of its own bunch crossing counter, a framing bit to indicate when the bits being sent are the LSBs, a signal to indicate whether 8-bit or 10-bit frames are being sent, and a parity bit for error detection. Each ADF channel is also equipped with a 16-bit pseudo-random generator and a constant value register that can drive its serial output when programmed to do so. The random traffic pattern generator is intended to test the communication link to the TABs independently of the digital filter. The constant register value can be used to debug the system (to know which channel goes where) or in normal data taking mode, to turn-off a noisy channel or switch off a particular digital filter (although this can also be accomplished by filling the E_T look-up table with the desired value).

Several possible implementations of the serial link have been proposed. These are discussed in more detail in section 4.7.

4.7.5.5 Clock and control signals

The ADF system receives from the Serial Command Link (SCL) interface the following signals:

- The 7.57 MHz beam-crossing clock. This clock is the only source of timing for the ADF system. All other clocks (for the ADC's, the logic and the output links) are derived from the 7.57 MHz beam-crossing clock.
- The initialize geographic section flag. This is used to initialize the ADF system. The ADF system returns an acknowledge signal when the initialization process is completed.
- The Turn Marker. Each ADF card increments locally a Beam Crossing counter and a Turn counter. These counters are intended to check that all ADF cards are properly synchronized.
- The L1 accept signal.
- One L1 qualifier. This signal is used to indicate that a special action has to be performed by the ADF cards for the current level 1 accept. A typical action is to freeze the content of all history buffers until these have been read-out via the slow control interface.

The ADF system returns to the trigger framework the following signals via the SCL interface:

- A initialize acknowledge signal.
- A L1 busy and a L1 error signal.

In addition to signals coming from the SCL, a number of signals need to be distributed to all ADF cards simultaneously. The timing distribution card is used to fanout these signals.

4.7.5.6 Slow Control Interface

Each ADF card includes a slave A24/D16 VME interface (Cypress CY7C964A and CY7C960A) connected to a standard J1 VME connector. This VME path is used to download FPGA configuration, to control each ADF card (adjust zero-offset DAC, program digital filter coefficients, download lookup tables…), to run tests programs, and to read-back the raw ADF samples that are captured following a monitoring request or following a self-trigger when the ADF card runs in standalone data acquisition mode.

4.7.5.7 Sources of trigger and modes of operation

The ADF system can respond to several sources of trigger and reacts to each of them in different ways. Each ADF channel can issue a self-trigger when the raw sample from the ADC considered is above a programmable threshold. Self-trigger signals are OR-ed together within each ADF board and are OR-wired at the level of each ADF crate. The self-trigger signal of each crate is relayed to the ADF Timing fanout card that dispatches this information to all ADF cards in a fully synchronous way. Upon reception of a self-trigger signal, each ADF card can freeze the content of its history buffers. These buffers can be read-out via the slow control interface before a synchronous order to resume recording is distributed. This mode of operation allows capturing pulses of various amplitudes

in selected channels, for signal shape characterization and for the determination of digital filter parameters.

Software triggers can be generated via the slow-control interface. These types of triggers are useful to capture complete views of the ADF inputs, intermediate results and final results on random beam crossings. This mode of operation can be used for system debugging in the absence of another source of trigger. It can also be used while the system is in operation, to verify that the ADF system does not mask any signal of significant amplitude. Technically, software triggers are generated by a write operation in the appropriate register of an ADF card. The timing fanout card relays the information in order to dispatch it to all ADF cards synchronously.

The principal source of triggers in normal data-taking mode is the Level-1 trigger. The Level-1 trigger information is distributed to all ADF cards from DØ Serial Control Link (SCL) using the ADF timing fanout card. A L1 Qualifier is also distributed along with each L1 accept signal. This L1 Qualifier is intended to indicate that the Trigger Control Computer (TCC) wishes to perform a monitoring operation for that L1 accepted event. For L1 accepted events with the L1 qualifier asserted, the ADF system freezes all its history buffer memories. The TCC reads out from each ADF card the relevant data and instructs the ADF system to resume raw samples and results recording. Note that while recording operations are suspended, the digital filter and sub-sequent logic continues to operate. Hence the monitoring operation does not introduce any dead time for data taking. Obviously a second monitoring request cannot be issued to the ADF system if it is already in the mode where data recording is suspended.

The ADF system can also be programmed to send raw samples to the TAB system following a Level 1 accept. When that functionality is disabled, L1 triggers are ignored by the ADF system (only L1 triggers with the L1 qualifier asserted are susceptible to be taken into account). When the function for sending raw samples to the TAB system is enabled, the following operations take place in the ADF upon reception of a L1 trigger:

- The content of all history buffer memories is frozen.
- Copies of the current beam crossing and turn number counters are latched in parallel-load serial-out registers
- A burst counter is loaded with the number of raw samples to be sent.
- The stream serialized to the TABs is switched from the output of the filter to the history buffer and a bit is set to indicate to the TABs that 10 bit frames are sent instead of 8 bit frames.
- Beam crossing number, turn number, and raw samples are sent.
- This operation repeats until the burst counter reaches 0.

When this operation is completed, the serial stream is switched back to the output of the filter and data recording in history buffers resumes automatically if

the event did not have the event qualifier asserted. In that later case, the ADF will send raw samples to the TABs and will wait for the software to resume recording. Sending raw samples after a Level 1 trigger to the TABs introduces a dead time in data taking proportional to the amount of data that is read-out. Assuming that a time window of 5 beam-crossings around that of interest is sent, 20 samples per channel are transmitted. The corresponding dead time is ~25 BC, i.e. 3.3 us. It should also be mentioned that the history buffer of each ADF channel is used both for read-out via the slow control and via the TABs. Following a qualified L1 trigger or a software trigger, raw sample recording is suspended until it is resumed by software. Consequently, for all the Level 1 triggers that occur during the corresponding period, none of the raw samples are captured and no raw data can be sent to the TABs (but filtered samples continue to flow).

In the ADF system, each source of trigger (self trigger, software trigger, L1 trigger and L1 qualifier) can be individually enabled/disabled.

4.7.5.8 Calibration, Monitoring and Data Acquisition

Calibration comprises two steps: coefficient determination and gain scaling. The first step consists in recording for each channel series of pulses in order to have sets of typical pulses shape. At this level, there is no need to correlate measurements with the energy that is effectively measured by the precision readout. Once filter coefficients have been determined, these need to be scaled to give the correct transverse energy value. The exact mechanism to correlate the raw ADC data that will be captured via the TCC to the output of the precision readout is still being discussed. In calibration mode, each ADF channel is programmed to record all raw samples until the occurrence of a self-trigger or software trigger.

Monitoring is done by capturing all ADC data on random beam crossings and for a fraction of L1 accept. This allows verifying that interesting events are not missed and that digital filters operate properly. Each ADF card includes sufficient buffering to keep all ADC raw samples and intermediate results during L1 latency (but no buffering beyond L1 latency is incorporated). No link between the ADF cards and the rest of DØ data acquisition exists. The only fast data path of the ADF cards is the links to the TAB. Following a L1 accept, the ADC boards do not re-send to the TABs the corresponding 2560 filtered energy values, but can optionally send the series of raw ADC values that correspond to the event that caused the trigger. Both raw samples and filtered energy values are therefore made available to the TAB system (with some restrictions as previously described).

4.7.5.9 ADF board implementation

Each ADF card is a 6U VME card that incorporates 32 channels. All analog and digital I/O connectors are placed on the rear side of the card. Analog inputs are placed on J2 connector, while digital outputs (cables connected to the TABs) are assigned to J0 connector. One ADF card per crate is also connected to the timing distribution card; the corresponding cable goes to pins on J0.

The analog circuitry is placed close to P2 connector, while FPGA's, the VME interface and Channel Link serializers are placed on the top half of the card. All the logic for the digital filters, buffer memories, look-up tables, etc, is implemented in Xilinx Virtex 2 FPGA's. This family was selected because it offers fast dedicated multipliers, large banks of dual ported SRAM, clock synthesizers, very compact shift registers as well as other attractive features. A 500K-gate chip accommodates eight channels; each ADF card included 4 such devices. All the internal logic is clocked at BC x 8 (60.6 MHz). Place and route software give a design running at 64 MHz for the slowest device speed grade (-4). Pin utilization (456 pin BGA model) is ~70, only 50% of the 32 internal multipliers and memory banks are used, but ~80% of the internal logic cells are utilized. A 250K-gate device does not comprise enough logic to accommodate 8 channels; a 1M-gate device does not have a sufficient number of I/O pins to implement 16 channels. Larger gate-count devices could have been employed, but it was found that these are less cost-effective than 4 500K-gate devices.

All the logic within the FPGAs is described in VHDL; there are ~70 source and test-bench modules in total.

4.7.5.10 ADF Crates and controller interface

There are 1280 trigger towers and each trigger tower comprises an EM and an HAD channel. There are 2560 analog channels in total. As each ADF card accommodates 32 channels, there are 80 ADF cards in total. These are housed in 4 fully populated 21-slot crates. Each crate contains 20 ADF cards and a VME Interconnect to make the interface to the Trigger Control Computer (TCC). Standard VME64x backplanes are used (5 rows connectors) but ADF cards only support A24/D16 slave transfers (i.e. 3U VME). The reasons that lead to the choice of a VME64x backplane are: the use of J0 to connect cables, pins for geographical addressing, lines for 3.3V power supply, the utilization of (normally reserved) bussed pins to fanout various clock and control signals within each crate, etc. A custom 20-slot wide passive backplane is placed at the back of the crate to make the transition between the 320 cables and connectors that bring the analog signals and the pins of RJ2 connectors. Because access to connectors at the back of the crate is critical in this design, the power supply of the crate needs to be placed below the card enclosure. The crate that was selected is Wiener model VME 6023 9U Plenum Bin.

4.7.6 Timing signals fanout card

This card is in charge of distributing to the ADF system all synchronous timing and control signals. It also distributes some intra ADF system global signals. A SCL mezzanine interface is plugged on the timing fanout card to receive and de-multiplex the signals transported by the DØ general synchronization network. The signals of interest have been previously mentioned. The only source of synchronization for the ADF system is the 7.57 MHz beam crossing clock. It must be guaranteed that all the results delivered by the ADF system to the TABs correspond effectively to the same beam crossing. Because of differences in the length of cables bringing the analog signals to the

ADF cards, there is a spread in time (~60 ns) between the peak observed on the earliest tower and that produced on the latest tower for any given beam crossing. In order to guarantee that a given beam crossing produces analog pulses that are seen by the ADF system as occurring during the same period of the BC clock, the following synchronization scheme is devised. Master timing signals are elaborated on the timing distribution card by inserting programmable delay lines on the relevant signals delivered by the SCL. All ADF cards receive a copy of these master timing signals via paths of equal length.

The timing card has 5 identical outputs links: 4 of these are connected to one ADF card within each ADF crate by a cable plugged in RJ0. The spare output is intended for test and debugging. The type of cable used is similar to these linking the ADF cards to the TABs. One of the cable wires is hard-wired to a know logic level on the timing card, while a resistor drives the corresponding pin to the opposite logic level on each ADF card. An ADF card detects the presence of the timing signals cable by sensing the logic level of the pin corresponding to the wire used for signaling. If the cable is present, that ADF card distributes the timing signals via the backplane bus to all the ADF cards within that crate (including itself). If no cable is detected, the ADF card places in tri-state all the outputs connected to the timing distribution lines. Any ADF card within a crate can act as a distributor of signals for that crate. However, in order to minimize ADF board-to-board signal skew, a card located close to the middle of the crate is preferably assigned that role of signal distributor. Given the differences in trace length, it is expected that the synchronization skew between any two ADF cards within the system is less than ~4 ns.

The timing card is also in charge of distributing to all ADF cards synchronous signals generated within the ADF system itself or issued by the TCC. These include for example, self-triggers and software triggers, the order to suspend and resume operation, the order to resume recording in the history buffers, etc. The different control commands are issued by the TCC to one of the ADF card that is a distributor of signals within a crate. The commands are transmitted to the timing card that makes the fanout to all ADF cards.

Because the number of signals to be exchanged between an ADF crate and the timing distribution card exceeds the number of pairs of wires available, a simple bit serial protocol is used to time multiplex the corresponding information over 2 half-duplex channels: a pair of wires transports data going upstream from the ADF to the synchronization card, a second pair transports data flowing in the opposite direction.

The timing distribution card is designed to avoid the need of a slow control interface. Nonetheless, several internal registers can be configured by the TCC through the ADF system by sending specific commands over any of the serial links going upstream from ADF crates to the timing card. Because all ADF crates are fully populated, there is no slot available at that level to accommodate the timing card. One option is to place it at the back of the controller in one of the ADF crates. The alternative is to house the ADF timing card in the TAB crate (the card only needs power and cooling).

4.7.7 Analog signal splitters

These cards are intended to test ADF card prototypes in the Run IIa setup without interfering with the CTFE system that is in operation. The analog splitter card receives the differential signals delivered by a BLS and makes two identical copies of these signals. One output is connected to a CTFE, the other output is available to an ADF. Because a split of the differential signals with a passive device would have lead to a reduction of the amplitude of the signals delivered to the CTFE and ADF, an active analog splitter is designed. Each channel of the splitter is composed of a fully-differential amplifier of gain $-\sqrt{2}$ driving a parallel branch of 2 fully-differential amplifiers of gain $-\sqrt{2}$. Hence, when both the input and the output are connected to properly terminated cards, the global gain of the splitter is unity. In order to reject common-mode noise as much as possible, fully differential operational amplifiers (THS4141) are used. All channels have a fixed gain; an optional capacitor can be soldered to limit the bandwidth individually. Each splitter card comprises 8 channels, i.e. 4 trigger towers. Connector spacing matches the pitch of the CTFE card to ease connections. Because only +5V and –5.2V are available within the crate housing the CTFEs while a swing from 0 to +5.5V of the analog signals is needed, an on-board DC/DC converter is used on each splitter card to generate locally +12V and –12V power supplies.

4.7.8 ADF to TAB Data Transfer

As discussed above, each of the 80 ADF cards will send three copies of data from 16 EM TTs and 16 H TTs to the TABs bit-serially in 8-bit words. Data transmission between the ADFs and the TABs will be accomplished using LVDS links. This solution is the simplest and most cost effective way of transmitting the large amount of data in the system. Two options were considered to implement the LVDS technology: solutions with LVDS drivers and receivers implemented in FPGAs or the use of LVDS driver/receiver chipsets. Both schemes are acceptable at the level of the ADC cards. The FPGA solution offers the advantages of elegance, flexibility and lower component count but requires more coordinated engineering than a ready-to-use chipset. Constraints at the TAB level favor the Channel Link solution and this is the solution that has been chosen for the baseline.

4.7.8.1 LVDS Serializers and Deserializers

The ADF cards to TAB links will be based on 48-bit Channel Link chipset[7] DS90CR483/484, which multiplexes its 48 parallel inputs to 8 data pairs. A total of 32 of the 48 inputs mapped to 6 output data pairs will be used to transmit the 16 EM and 16 H TTs on each ADF. Input data is clocked into the chip at 8xBC (60.6 MHz) while the output pairs run at 364 MHz. The format of the data transported on the links, including a list of possible error detection and synchronization characters is given in Table 14.

[7] National Semiconductor: Channel Link Chipset – DS90CR484, http://www.national.com/pf/DS/DS90CR484.html.

Table 14: Data transport protocol for LVDS links between ADFs and TABs.

Pins	Signal	Input Data Streams (0-31)
	Differential Data	0,1,2,3,4,5
	Differential Data	8,9,10,11,12,13
	Differential Data	16,17,18,19,20,21
	Differential Data	6,7,14,15,22,23
	Differential Data	24,25,26,27,28,29
	Differential Data	30,31
	Differential Clock	
	Framing bit	
	Select 8/10-bit frames	
	Parity	
	Grounds	

4.7.8.2 Cables

The cabling diagram between the 80 ADFs and 8 TABs is shown in Figure 58. Each ADF sends out 3 cables carrying identical copies of its data. Each TAB accepts 30 ADF cables giving it more than the 40×9 TTs of data in $\eta\times\phi$ that it requires to find RoIs over the full eta range and 4 phi slices.

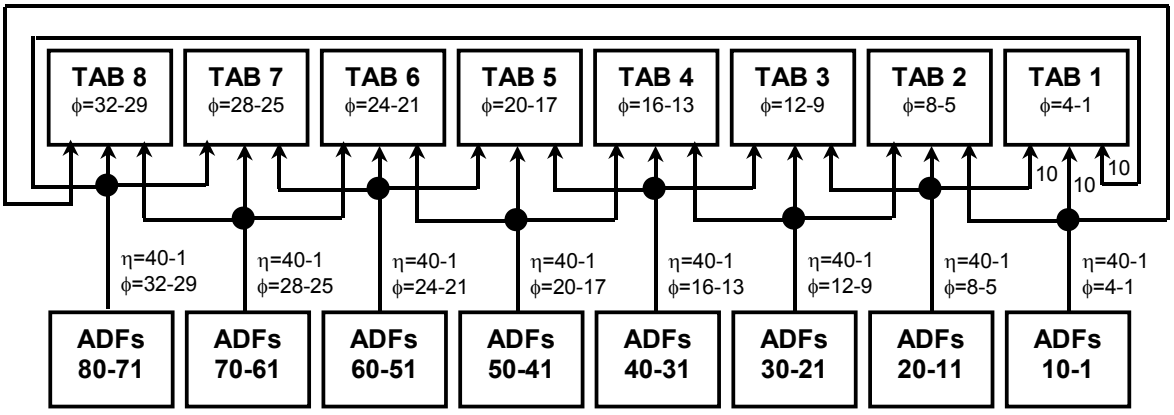

Figure 58: Diagram of cabling between the ADFs and the TABs.

4.7.9 TAB implementation

The Trigger Algorithm Boards (TAB) are the main processing component of the L1 calorimeter trigger. They perform the following tasks for a specific $\eta \times \phi$ region of the calorimeter.

- Find EM, Jet and tau clusters and estimate their E_T using sliding windows algorithms.
- Send a list of EM, Jet and tau clusters passing several different E_T thresholds to the GAB.
- Send a list of EM and Jet clusters over threshold to the Cal-Track Match system.
- Calculate a scalar E_T sum and E_x and E_y sums for all TTs in the TAB's region for use in $E_{T,total}$ and Mp_T.
- Send the E_T, E_x and E_y sums to the GAB.
- Format and buffer data for use in L2 and L3 and send this out on receipt of L1 and L2 trigger accepts.
- Allow insertion of fake data, by computer, directly after the inputs to test system functionality.

4.7.9.1 Overall TAB Architecture

A block diagram of the TAB is given in Figure 59. The architecture is driven by several global considerations.

- Our target algorithms use 2×2 windows and 5×5 declustering regions.
- ADF design is simplified if each such board is identical. Practically this means that the same number of output cables must be sent from each board.
- A large number of TT signals must be sent to a given TAB to find E_T cluster local maxima
- Latency must be kept to a minimum
- Small, low-speed FPGAs should be used to minimize costs.

These considerations have led us to a design in which each TAB checks for local maxima in a region covering the full η region by 4 bins in ϕ. This means that each TAB will output information on 31×4 possible local maxima. Note that there are only 31 possibilities in η because of the four TTs devoted to the ICR and the five TTs at the edges of the calorimeter for which not enough data exists to decluster.

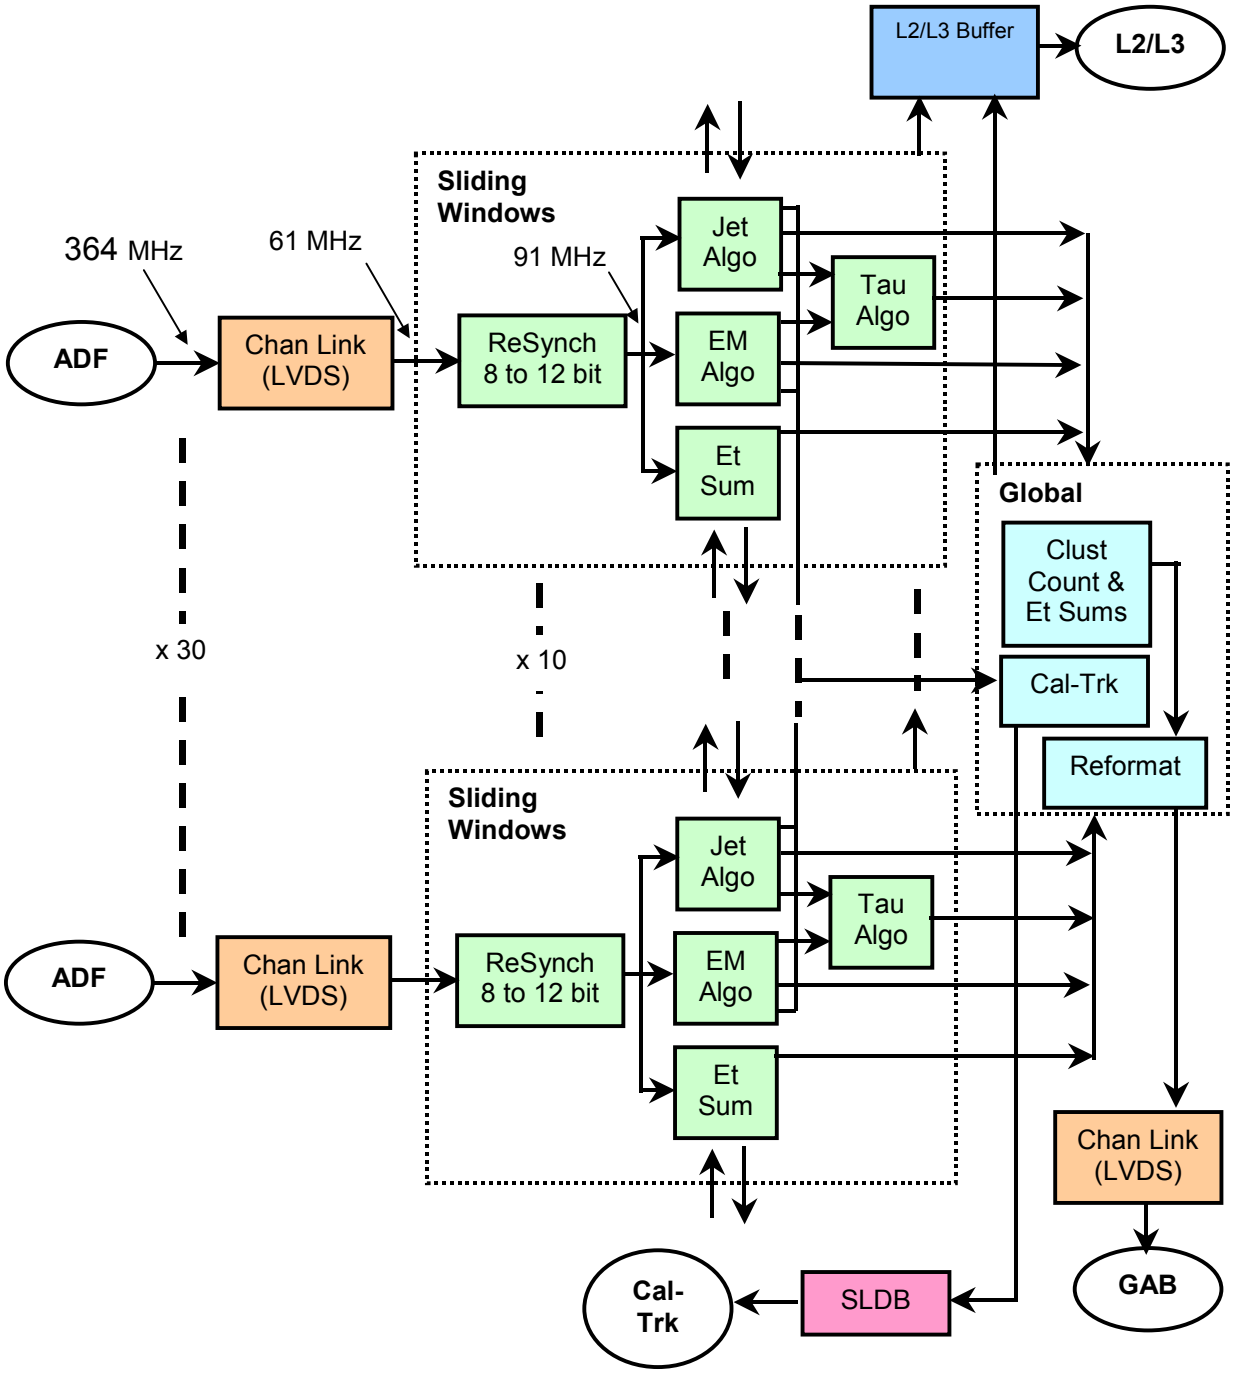

Figure 59: Block diagram of the TAB.

The functional elements of the TABs are the following.

1. LVDS receivers for input signals from the ADFs (more details are given in the next section). These produce as output 32 streams of 8-bit serial data corresponding to the EM and H E_T's for each bunch crossing (BC) from the ADF at 60.6 MHz (8xBC).

2. Resynchronizers, which align the serial data streams and retransmit them as 12-bit serial streams, padded with zeroes, at 90.9 MHz (12xBC). The 12-bit range of these outputs is necessary to deal with carries in the serial adders described below. The resynchronizers also fanout the signals for use in several sliding windows chips.
3. Pre-summers that add ICR energies to the calorimeter EM or H TTs from the region to which they belong.
4. Sliding Windows blocks that implement the sliding windows algorithms for both EM, Jet and tau clusters for a sub-set of the data considered by the board and also perform the first step of global summing.
5. A Global FPGA that gathers information from the Sliding windows FPGAs, constructs the list of RoIs passing EM and Jet thresholds, does another step in the global summing chain and sends out information from this TAB to the GAB. This chip also prepares data for transmission to the Cal-Track Match system.
6. An LVDS transmitter that receives data from the Global FPGA at 90.9 MHz and sends it to the GAB.
7. At various stages in the processing data is buffered for transmission to L2 on the receipt of an L1 accept and for transmission to L3 on an L2 accept.

Preliminary versions of firmware have been written for all relevant elements of this chain and candidate devices have been chosen.

4.7.9.2 Serial Arithmetic Building Blocks

The sliding windows algorithm is basically a set of sums and compares. These operations can be performed efficiently in a bit-serial manner that minimizes FPGA resources required, particularly the number of input data pins. This serial structure also meshes well with the serialized data arriving from the ADFs. The two basic arithmetic/logic elements required for the sliding windows algorithm are the addition of two numbers and the compare of two numbers. Nearly all elements of the algorithm can be constructed from trees of these primitives. Diagrams of a bit-serial adder for two data lines and for a bit-serial comparator, as implemented in first versions of the sliding windows firmware, are shown in Figure 60 and Figure 61.

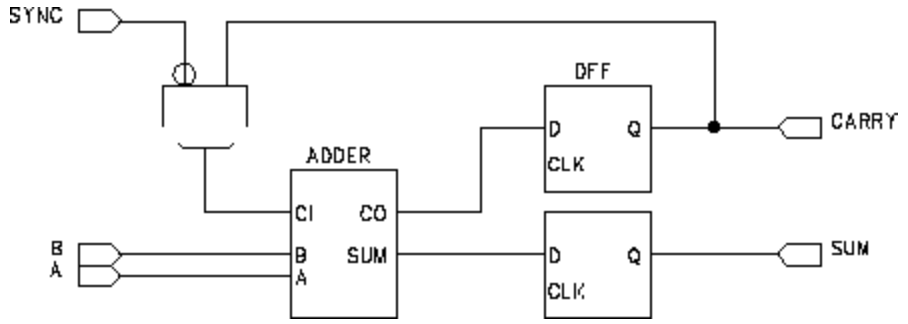

Figure 60 : Serial adder for data A and B. SYNC is a signal that separates one serial word from the next.

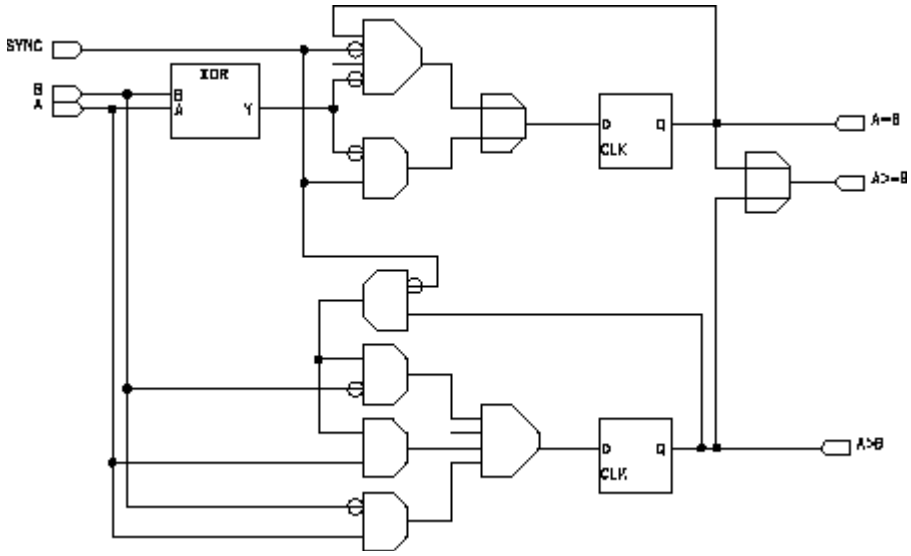

Figure 61: Serial comparator for data A and B. Outputs "A>B" and "A>=B" are both required to unambiguously assign local maxima.

4.7.9.3 Sliding Windows Algorithms Implementation

The EM, jet and tau cluster finding algorithms described in Section 4.6 have been implemented on a set of FPGAs on the TAB. Each FPGA finds EM, jet and tau clusters in a sub-set of the data considered by the TAB. Data sharing is discussed in more detail below.

Data distribution to the sliding windows chips

Data distribution within the TABs is complicated because of the large amount of data sharing required by the sliding windows algorithm as well as by the desire to send ICR energies to all chips that need them. Additionally, the FPGAs in which the algorithms are implemented have only four external clocks. For this reason each chip can take direct input data from only three ADF cables (each TAB receives 30 such cables). This sets the number of sliding windows chips at 10 per TAB. Each chip must pass along some of its input data to its neighbor chips so that each chips gets a 9×9 region of TTs, allowing it to construct (up to) 4×4 local maxima.

An example of the intra-TAB data paths for TAB-2 (dealing with ϕ's 5-8) is shown in Figure 62. The diagram shows cable inputs to each chip, the data it shares with its neighbors, the distribution of ICR data (in η's 1,2 and 39,40) to all the chips that need it and the chip's output local maxima.

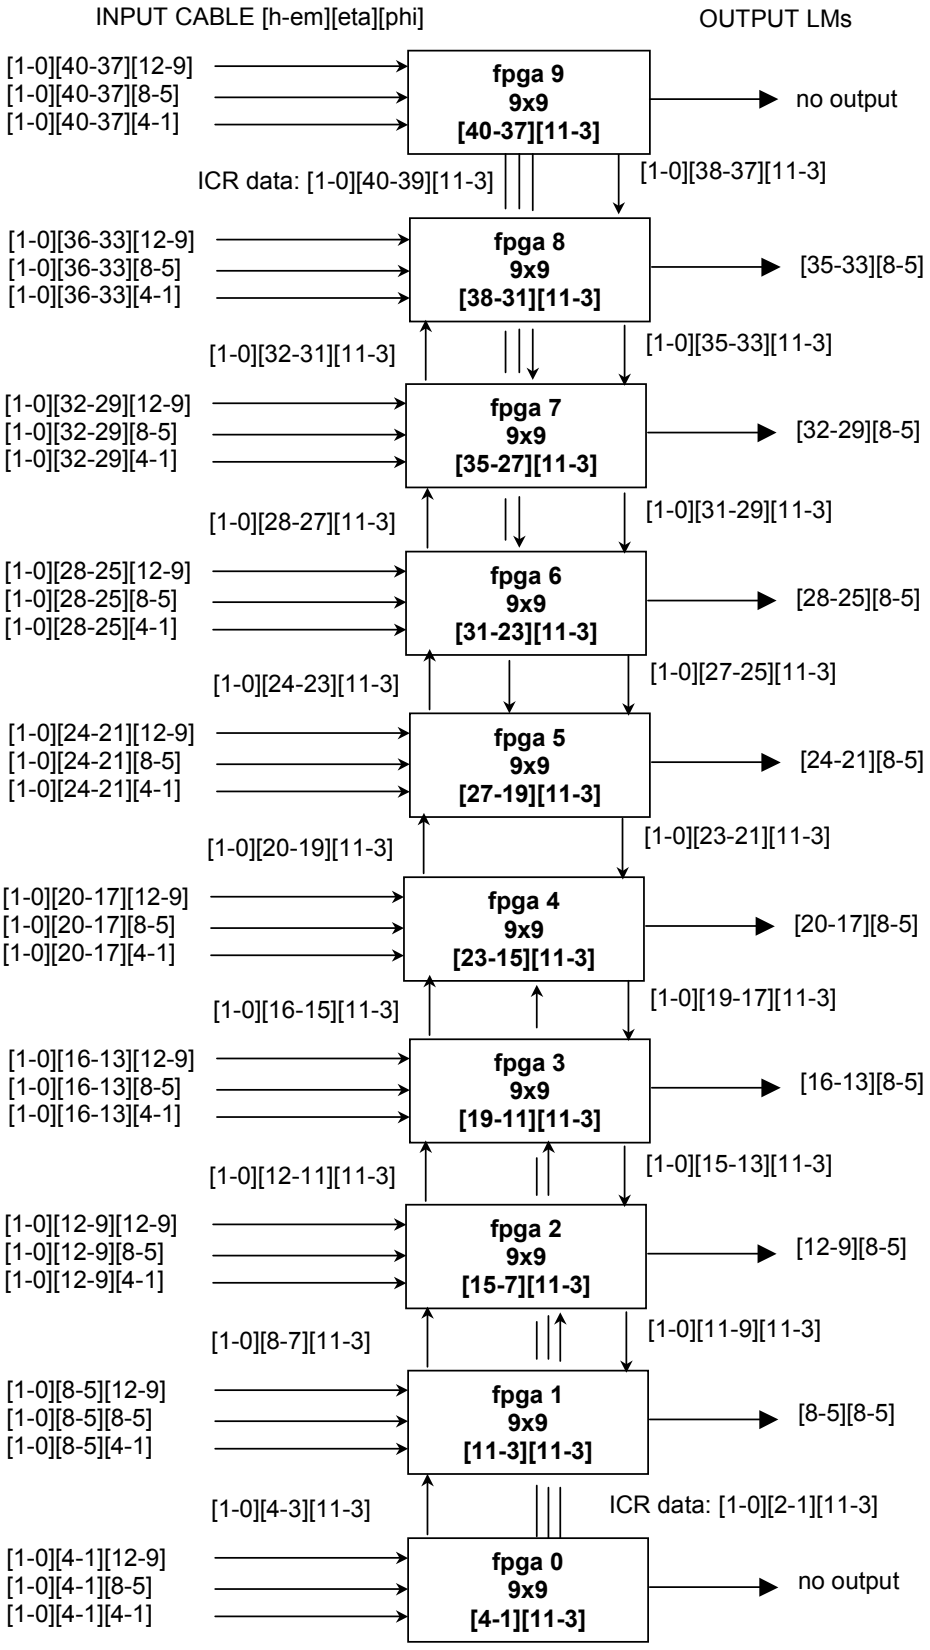

Figure 62: Data sharing within a TAB.

EM and jet cluster algorithms

As an example of the firmware design, a schematic of the EM algorithm (without the local maximum finding) is shown in Figure 63. It includes the RoI sum, EM isolation sum and Had isolation sum, a comparison of the EM isolation sum with a downloaded threshold and a comparison of the Had isolation sum with Roi-sum/8 (shifted by 3 bits). Also included are threshold comparisons for the RoI for five thresholds. The local maximum finding schematic is shown in Figure 64 where the RoI labeled “E[1][1]” is the candidate RoI.

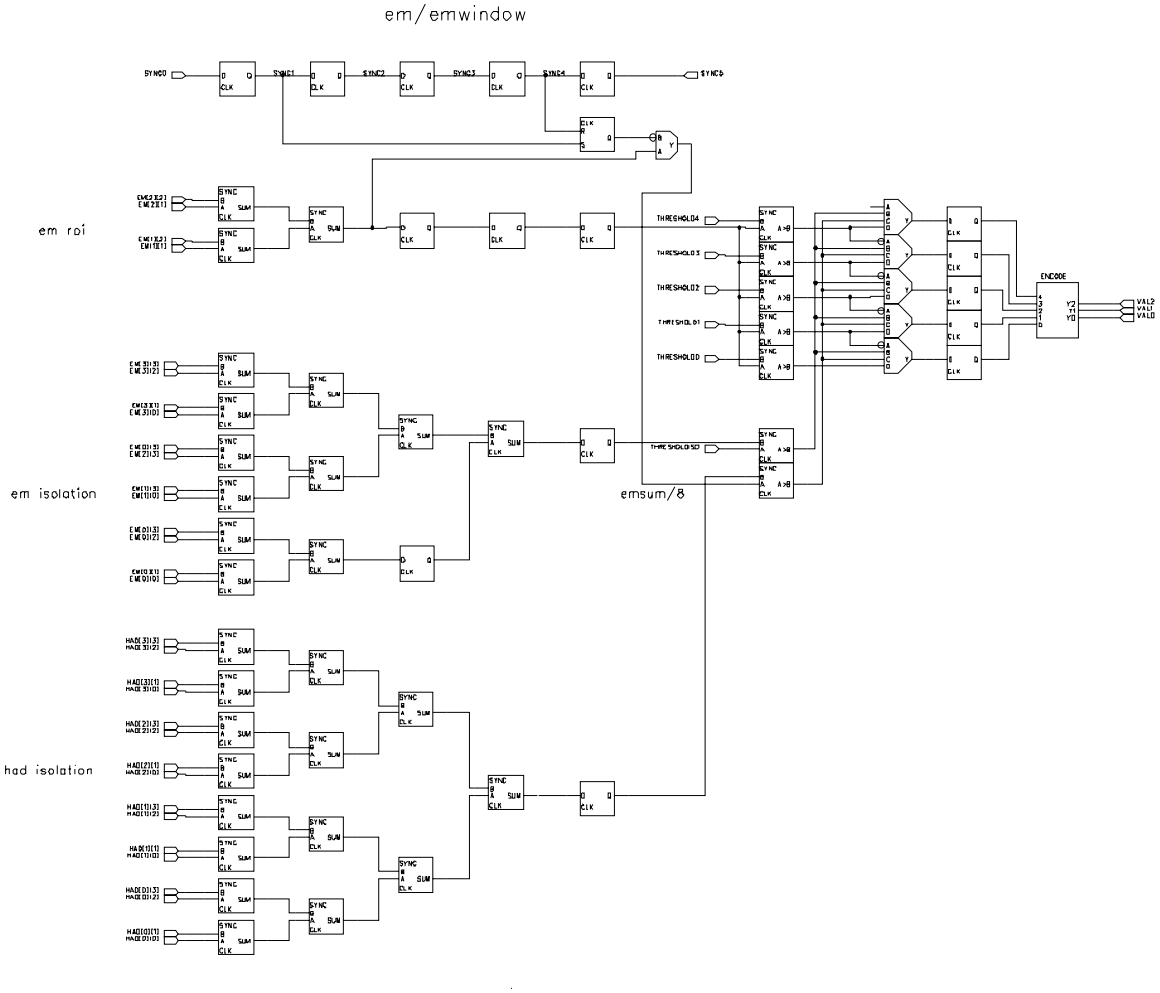

Figure 63: Schematic for the EM sliding windows algorithm.

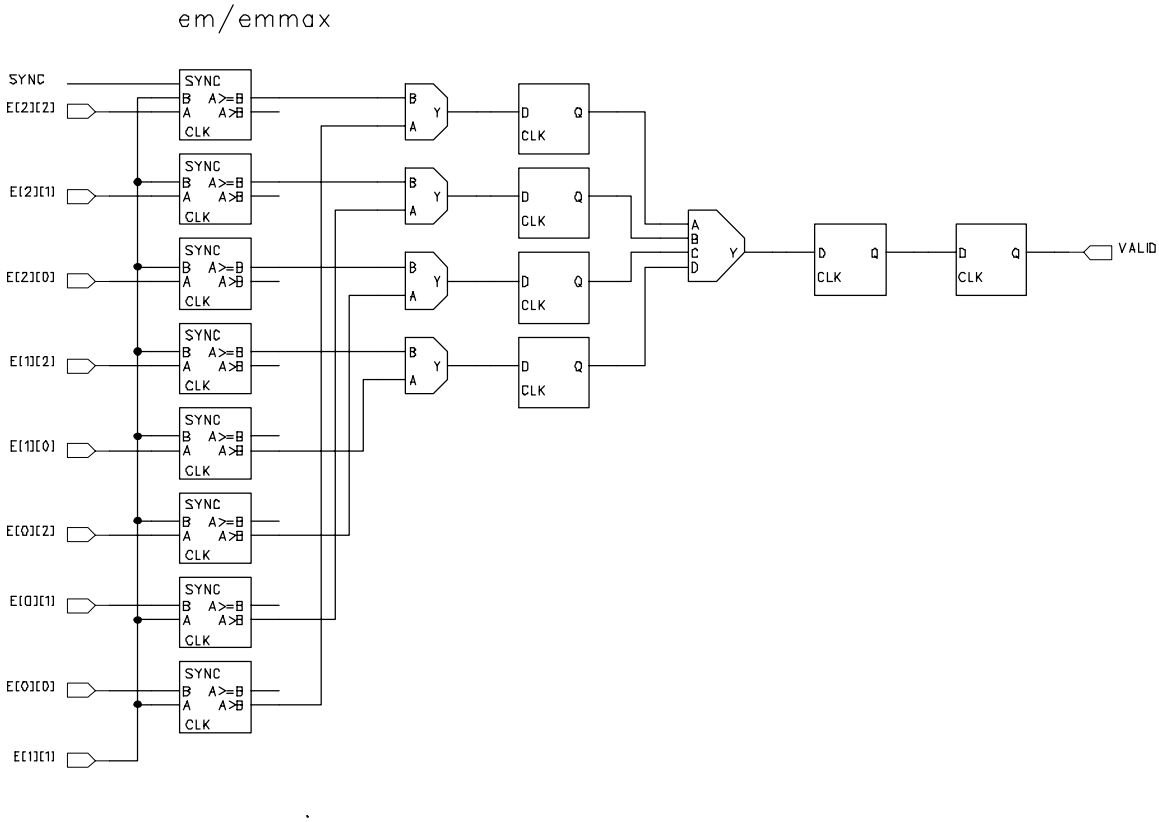

Figure 64: Schematic of local maximum-finding algorithm for the EM sliding window.

Tau Algorithm

The tau algorithm is run after jet declustering has been performed for those windows which are found to be local maxima. It uses the 2×2 and 4×4 sums calculated as part of the jet algorithm. The ratio cut is implemented as a small Look-Up-Table.

Threshold Comparisons

The E_T's of EM and Jet clusters that are found to be local maxima by the sliding windows algorithm are compared with two sets of up to seven E_T thresholds, which can be downloaded at initialization. The results of these comparisons are numbers 1-7 encoding the highest threshold passed. The Sliding Windows chip then sends out information about all RoIs that it considers, as 3 bits for each RoI packed with 4 RoIs a in 12-bit word. The 3 bits for a given RoI are either 0, if the RoI is not a local maximum, or 1-7, corresponding to the highest threshold it has passed if it is a local maximum.

Clusters found by the tau algorithm, for which one threshold is foreseen, are transmitted as a 3 bit encoded threshold as for the jets and EM clusters.

Table 15: Cluster encoding for data transmitted from sliding windows chip to global chip.

3-bit Value	Meaning
0	RoI is not a local maximum
1-7	RoI is a local maximum Value = index of highest Et threshold passed

E_T Sums

Preliminary sums for use in calculating $E_{T,total}$ and Mp_T are also computed on the Sliding Windows chips. Each chip sums all TT EM and EM+H E_T's over all η for each of the four ϕ's considered in the chip.

Clusters for the Cal-Trk Match system

For use by the Cal-Trk match system, each sliding windows chip constructs a count over all η of the EM and jet clusters passing each of four thresholds. In order to save the time, this information is sent on to the Global chip in parallel before the main threshold comparisons.

Sliding Windows Chip Outputs

The outputs of the Sliding Windows chips sent to the Global chip are then a set of 12-bit words sent serially at 90.9 MHz as shown in Table 16 and the 32 lines for fast Cal-Trk match data, corresponding to four thresholds for EM and jet clusters for each of the four ϕ's dealt with by the chip (Table 17).

Table 16: Sliding windows chip cluster and sum output data format.

No.	Type	11	10	09	08	07	06	05	04	03	02	01	00
1	EM	η=4, ϕ=1			η=3, ϕ=1			η=2, ϕ=1			η=1, ϕ=1		
2	Jet	4,1			3,1			2,1			1,1		
3	Tau	4,1			3,1			2,1			1,1		
4	EM	4,2			3,2			2,2			1,2		
5	Jet	4,2			3,2			2,2			1,2		
6	Tau	4,2			3,2			2,2			1,2		
7	EM	4,3			3,3			2,3			1,3		
8	Jet	4,3			3,3			2,3			1,3		
9	Tau	4,3			3,3			2,3			1,3		
10	EM	4,4			3,4			2,4			1,4		
11	Jet	4,4			3,4			2,4			1,4		
12	Tau	4,4			3,4			2,4			1,4		
13	ΣE_T	EM E_T: sum over η; ϕ=1											
14	ΣE_T	EM+H E_T: sum over η; ϕ=1											
15	ΣE_T	EM E_T: sum over η; ϕ=2											
16	ΣE_T	EM+H E_T: sum over η; ϕ=2											
17	ΣE_T	EM E_T: sum over η; ϕ=3											
18	ΣE_T	EM+H E_T: sum over η; ϕ=3											
19	ΣE_T	EM E_T: sum over η; ϕ=4											
20	ΣE_T	EM+H E_T: sum over η; ϕ=4											

Table 17: Dedicated output for EM and jet clusters over threshold for Cal-Trk match.

No.	Type	11	10	09	08	07	06	05	04	03	02	01	00
1	EM	Count: ϕ=1, Thr=4			Count: ϕ=1, Thr=3			Count: ϕ=1, Thr=2			Count: ϕ=1, Thr=1		
2	Jet	Count: ϕ=1, Thr=4			Count: ϕ=1, Thr=3			Count: ϕ=1, Thr=2			Count: ϕ=1, Thr=1		
3	EM	Count: ϕ=2, Thr=4			Count: ϕ=2, Thr=3			Count: ϕ=2, Thr=2			Count: ϕ=2, Thr=1		
4	Jet	Count: ϕ=2, Thr=4			Count: ϕ=2, Thr=3			Count: ϕ=2, Thr=2			Count: ϕ=2, Thr=1		
5	EM	Count: ϕ=3, Thr=4			Count: ϕ=3, Thr=3			Count: ϕ=3, Thr=2			Count: ϕ=3, Thr=1		
6	Jet	Count: ϕ=3, Thr=4			Count: ϕ=3, Thr=3			Count: ϕ=3, Thr=2			Count: ϕ=3, Thr=1		
7	EM	Count: ϕ=4, Thr=4			Count: ϕ=4, Thr=3			Count: ϕ=4, Thr=2			Count: ϕ=4, Thr=1		
8	Jet	Count: ϕ=4, Thr=4			Count: ϕ=4, Thr=3			Count: ϕ=4, Thr=2			Count: ϕ=4, Thr=1		

4.7.9.4 Global FPGA

The Global FPGA receives information from the TABs outlined in Table 16 and Table 17. Its main job is to collect the encoded EM and Jet data for output and calculate E_x and E_y, derived from the E_T ϕ-sums bit-serially using x,y weights stored in ROM as a Look-Up-Table. It also counts the number of EM, jet and tau clusters over threshold.

The Global chip also provides data to the Cal-Track match system based on the fast Cal-Trk data lines. OR's of the threshold bits are constructed over all η and counts are made of the number of EM and jet clusters passing one of the thresholds. This information is then packed into 8-bit words – one for each EM and jet cluster considered by the TAB. The bit definition of these words is given in Table 18.

Table 18: Format of 8-bit Cal-Trk match word sent from Global chip. One such word is sent for each ϕ for EM and Jet values.

07	06	05	04	03	02	01	00
				Threshold bits ORed over η			
cluster count – Thr n				Thr 4	Thr 3	Thr 2	Thr 1

4.7.9.5 Output to the GAB

The Global FPGA on each TAB produces the following types of data.

1. Counts of EM and jet clusters over each of the seven thresholds.
2. A count of the number of tau clusters over threshold.
3. EM and EM+H E_T sums for the region considered by the TAB.
4. EM+H E_x and E_y sums for the region considered by the TAB.

Each TAB considers $31 \times 4 = 124$ possible local maxima. However, since the 2,1,1 algorithm requires that local maxima be separated by at least two TTs in η and ϕ, only 31 EM and 31 jet clusters are allowed. Thus, cluster counts can be safely stored as 6-bit words.

The TABs will send their data out to the GAB serially using the same LVDS system as used for the ADF to TAB transmission. In this case, the TABs will send data as 12-bit words using the format given in Table 19.

Table 19: Data format for TAB to GAB transfer.

No.	11	10	09	08	07	06	05	04	03	02	01	00
1	Jet Count: Threshold=1						EM count: Threshold=1					
...	...						...					
7	Jet Count: Threshold=7						EM count: Threshold=7					
8							Tau count					
9	EM E_T: sum over η, ϕ=1											
10	EM+H E_T: sum over η, ϕ=1											
...	...											
15	EM E_T: sum over η, ϕ=4											
16	EM+H E_T: sum over η, ϕ=4											
17	EM+H E_x											
18	EM+H E_y											

4.7.9.6 Output to Cal-Trk match system

Data is sent to the Cal-Trk system using an existing L1Muon Serial Link Daughterboard[8]. The Cal-Trk system expects 16-bit words at a frequency of 53 MHz. A maximum of 7 such words can be sent in one B.C. for a total bit-count on a single cable of 112 bits per B.C. The number of cables from L1Cal to the Cal-Trk system is limited to approximately 12, which is safely above the 8 sent by the baseline system.

During periods with no data (SYNC_GAP, etc.) the Cal-Trk receiver expects to receive NULL (K28.5) characters as data words.

Several control bits are required by the Cal-Trk system[8]. The two of most relevance to L1Cal transmission are:

- Enable (E): high = data enabled
- Parity_Enable (P): high = data is longitudinal parity

The Global FPGA formats data according to the specifications of the L1 Muon SLDB. Its output will have the form given in Table 20.

[8] Serial Link Daughter Board Specification, http://hound.physics.arizona.edu/l1mu/l1mu.htm.

Table 20: Format for data from a TAB to the Cal-Trk match system

No	P	E	15	14	13	12	11	10	09	08	07	06	05	04	03	02	01	00
…	0	0	NULL															
1	0	1	Jet: count + E_T: ϕ=1								EM: count + E_T: ϕ=1							
2	0	1	Jet: count + E_T: ϕ=2								EM: count + E_T: ϕ=2							
3	0	1	Jet: count + E_T: ϕ=3								EM: count + E_T: ϕ=3							
4	0	1	Jet: count + E_T: ϕ=4								EM: count + E_T: ϕ=4							
5	1	1	Longitudinal Parity															

4.7.10 GAB Implementation

The Global Algorithm Board (GAB) is the final step in the L1 calorimeter chain before the Trigger Framework. A list of the tasks that this board will perform is given below.

- Calculates $E_{T,total}$ and E_x, E_y from the TAB partial sums and transmits them to the Trigger Framework
- Constructs trigger terms from counts of clusters over threshold and global sums and sends these to the trigger framework.
- Receives timing and control information from the Trigger Framework over the Serial Command Link (SCL) via an SCL Receiver mezzanine card[9] (SCLR) and fans this out to the ADFs and TABs as necessary.
- Sends data to L2 and L3.

4.7.10.1 Output to the Trigger Framework

Trigger terms sent from the GAB to the Trigger Framework are still being explored. All quantities currently sent to the Trigger Framework by the L1 calorimeter trigger (see Section 4.3) will continue to be sent though.

4.7.10.2 Control and Timing Signals for the TAB-GAB System

The TABs need to receive global control signals such as a bunch crossing clock and event number. This information is available from the Trigger Framework via the Serial Command Link (SCL) using custom SCL Receiver mezzanine cards[9]. However, the TABs need only a small subset of the full SCL information. It is therefore more efficient to receive this information centrally in the GAB and then fan out the required sub-set to the TABs.

A list of SCL signals, their sources and their destinations within the system is given in Table 21.

[9] DØ Trigger Distribution System: Serial Command Link Receiver (SCLR), http://www-ese.fnal.gov/d0trig/default.htm.

Table 21: Signals from the SCL and their routing in the TAB-GAB system.

Data	Source	Destination	Comments
SCL_ACK	GAB (+TAB)	SCLR	in reset
SCL_READY	SCLR	GAB	
SCL_SYNCERROR	SCLR	GAB	
SCL_DATAERROR	SCLR	GAB	
CLK_53	GAB	TAB	main clock
CLK_7			
CURRENT_TURN[15..0]	GAB	TAB: L2/L3	needed for L2 header
CURRENT_BX[7..0]	GAB	TAB: L2/L3	needed for L2 header
FIRST_PERIOD			
BEAM_PERIOD			
SYNC_GAP	SCLR	GAB	for Framework data
COSMIC_GAP			
SPARE_PERIOD			
L1_PERIOD	GAB	TAB: L2/L3	L2 transmission / monitoring
L1_ACCEPT	GAB	TAB: L2/L3	L2 transmission / monitoring
L1_TURN[15..0]	GAB	TAB: L2/L3	L2 data error check
L1_BX[7..0]	GAB	TAB: L2/L3	L2 data error check
L1_QUAL[15..0]	GAB	TAB:L2/L3+Mon +Cal-Trk	L2/Mon control + Cal-Trk
L2_PERIOD			dealt with by VRB
L2_ACCEPT			dealt with by VRB
L2_REJECT			dealt with by VRB
INIT_SECTION	SCLR	GAB,TAB	initialize
L1_BUSY	GAB TAB	SCLR GAB	inputs backing up
L2_BUSY			dealt with by VRB
L1_ERROR	GAB TAB	SCLR GAB	GAB error TAB lost synch, etc
L2_ERROR			
INIT_ACK	GAB (+TAB)	SCLR	acknowledge init request
SYNC_LOST	GAB	SCLR	synch lost
SPARE_STATUS[1..0]			

4.7.11 TAB-GAB Crate

The TABs and GAB will be housed in one 9U VME crate. This crate will have standard VME connectors on P0, P1 and P2. A VME CPU will also be included in this crate for use in downloading, monitoring and testing.

4.7.12 Output to L2 and L3

The Run IIa L1 calorimeter trigger currently sends the same data to L2 and L3 using G-Link fiber optic cables. Ten such cables are optically split to provide data from each of the ten racks in the system, corresponding to an $\eta \times \phi$ region of 4x32 TTs. The data structure is given in Table 22. In the table, "seed masks" have a bit set if the corresponding TT has E_T above a run-time defined threshold. The generic requirements for L2 data, including header and trailer definitions are given on the L1CTT web pages[10].

Table 22: Run IIa L1 calorimeter data to L2 and L3.

Starting Byte #	# of Bytes	Data
1	12	L2 Header
13	16	EM TT seed mask
29	16	EM+H TT seed mask
45	128	EM TT E_T's
173	128	EM+H TT E_T's
301	4	L2 Trailer

A similar data set will be transmitted to L2 and L3 in the new system, with the TT seed masks being replaced by encoded cluster data. However, more information may be desirable. The necessity of passing along all of the TT E_T's will require that this L2/L3 data be sent from the individual TABs. Preliminary versions of L2/L3 data from the TABs and GAB are given in Table 23 and Table 24.

[10] L1CTT homepage – Current Protocols (v07-00): http://d0server1.fnal.gov/projects/VHDL/General/

Table 23: Preliminary format of data from TAB boards to L2 and L3.

No.	15	14	13	12	11	10	09	08	07	06	05	04	03	02	01	00
1	EM+H TT Et: eta=1, phi=1								EM TT Et: eta=1, phi=1							
2	EM+H TT Et: eta=2, phi=1								EM TT Et: eta=2, phi=1							
…	…								…							
40	EM+H TT Et: eta=40, phi=1								EM TT Et: eta=40, phi=1							
41	EM+H TT Et: eta=1, phi=2								EM TT Et: eta=1, phi=2							
…	…								…							
160	EM+H TT Et: eta=40, phi=4								EM TT Et: eta=40, phi=4							

Table 24: Preliminary format for data transfer from GAB to L2 and L3. Note that there are only 31 possible cluster positions in eta because of edge effects.

No.	15	14	13	12	11	10	09	08	07	06	05	04	03	02	01	00
1	Mask of EM clusters over L2 threshold: eta=5, phi=1-16															
2	Mask of EM clusters over L2 threshold: eta=5, phi=17-32															
3	Mask of EM+H clusters over L2 threshold: eta=5, phi=1-16															
4	Mask of EM+H clusters over L2 threshold: eta=5, phi=17-32															
…	…															
121	Mask of EM clusters over L2 threshold: eta=35, phi=1-16															
122	Mask of EM clusters over L2 threshold: eta=35, phi=17-32															
123	Mask of EM+H clusters over L2 threshold: eta=35, phi=1-16															
124	Mask of EM+H clusters over L2 threshold: eta=35, phi=17-32															
125	0	0	0	0	EM Sum Et											
126	0	0	0	0	EM+H Sum Et											
127	0	0	0	0	EM+H Sum Ex											
128	0	0	0	0	EM+H Sum Ey											

4.7.13 Latency of Output to the Cal-Track Match System

We have estimated the latency in sending data from the L1 calorimeter trigger to the Cal-Trk match system using simulations of firmware in the ADFs and TABs. See Section 4.7.9.6 for the details of the data transferred to the Cal-Trk boards. Our latency estimate is given in Table 25. The current implementation of the L1muon trigger, upon which the Cal-Trk system is based, requires data to arrive at its input no later than 1500 ns after the BC producing that data, we are clearly over budget. Thus, we need to increase the amount of time available for the Cal-Trk system to make its decision. This is set by the shortest pipeline depth over which a DØ sub-detector can hold its data waiting for

an L1 decision. The limiting pipelines will come from the muon central drift tubes and forward scintillators. (The silicon detector will have a new readout chip with a significantly deeper pipeline and the fiber readout will use the new TRIP readout chip that will replace the SIFT/SVX2.) We have determined that the muon pipelines can be lengthened by six 132 ns clock ticks without affecting the data quality. As can be seen from the table below, with this change adequate latency for the cal-track match system is achieved.

Table 25: Latency estimates for data transmission to the Cal-Trk match system from the L1 calorimeter trigger

Step	Δt (ns)	Time from BC (ns)	Comments
BC to ADF	650	650	transmission from calo to ADF inputs
Digitization	347	997	ADF
Digital Filter	594	1591	
Output to TAB	206	1797	
Resynch Inputs	132	1929	TAB
Input Data Checks	66	1995	
Jet/EM Algorithms	242	2237	
Construct Output	198	2435	
SLDB	90	2525	
Cal-Trk requirement		3363	muon pipeline depth increased
		2571	“as is” system

4.8 Summary & Conclusions

The high luminosity of Run IIb presents a significant challenge to the L1 calorimeter trigger. The L1 calorimeter trigger upgrade proposed in this section addresses these challenges.

We will need to be able to effectively identify calorimeter energy depositions with the correct bunch crossing – this is addressed by digital filtering techniques. Since almost 80% of the L1 rate is calorimeter based, the importance of sharpening the p_T cut (and thus reducing background rates) as well as the possibility of triggering on real objects such electromagnetic clusters and jets is clear, and being addressed by a “sliding window” technique.

The improvement in global variables such as missing E_T can also be improved with the addition of the energy from the ICR region at L1. The ability to do that has been provided in the present front-end electronics.

Finally, the additional power provided by current FPGA’s will allow the migration to L1 of more sophisticated algorithms and topological cuts presently available at L2.

This set of tools provided by the proposed L1 calorimeter trigger will allow us to make the optimal use of the Run IIb luminosity.

The hardware implementation of these improvements has been explored, leading to an overall architectural design. Preliminary detailed designs for several of the most important elements of the system have been made, with a full first-pass design anticipated by summer 2002.

5 Level 1 Calorimeter-Track Matching

5.1 Overview

The goal of the L1CalTrack trigger is to exploit matches in the ϕ position of tracks from the L1CTT trigger with that of EM and jet objects from the L1Cal trigger in order to reduce the L1 trigger rates of EM and track triggers. Information from the Central Preshower (CPS) and Forward Preshower (FPS) detectors is also used. Monte Carlo studies show that the improvement in the reported ϕ position of EM objects at the trigger level from 90º to 11.25º can reduce medium P_T electron triggers by a factor of 2-3. Additionally, large factors of rejection (10-70) can be achieved by matching track triggers with calorimeter towers of modest energy. This latter is important in triggering on hadronic tau decays such as in $H \rightarrow \tau^+ \tau^-$.

The implementation of the L1CalTrack trigger uses the **existing** L1Muo architecture with small modifications. This is sensible since the L1Muo trigger matches the ϕ position of tracks from the L1CTT trigger with that of muon objects derived using muon scintillation counter hits, a similar function to the L1CalTrack trigger. The huge advantage of this implementation is that the L1Muo trigger has been successfully running since the start of Run 2. Thus issues such as synchronization, buffering, outputs to L2 and L3, electronics testing, monitoring, power supplies, and rack infrastructure have proven, working solutions.

5.2 Simulation

5.2.1 Improving Calorimeter EM Rates Using the L1CTT

The L1CTT trigger is not yet operational on DØ hence we rely on Monte Carlo studies at present to estimate the gains of an L1CalTrack trigger. The simulation effort must be improved (by including CPS and FPS information for example) and cross-checked with collider data. Nevertheless, the existing Monte Carlo studies indicate that Run IIa electron and track trigger rates can be reduced by the addition of the L1Cal Track trigger. The reasons for this rejection are the improved ϕ granularity of EM and jet objects from L1Cal and the fact that the fake rates in the calorimeter and central fiber tracker are relatively uncorrelated.

This latter point is shown in the following studies that match EM objects from the calorimeter with tracks from the L1CTT. The calorimeter EM objects are found with different E_T thresholds. The L1CTT tracks have P_T > 1.5 GeV/c unless otherwise noted. A calorimeter–track match is defined by matching the calorimeter EM trigger towers ($\Delta\phi$=11.25º) with tracks from the three overlapping L1CTT track sectors (each $\Delta\phi$ = 4.5º). QCD jet events were used to simulate the background.

Results are shown in Table 26. The left-hand column gives the E_T threshold for the calorimeter EM objects. The denominator in the subsequent columns is the number of EM objects (trigger towers) exceeding each E_T threshold. The numerator is the number of EM object-track matches. Results at

two different luminosities and two different QCD regimes are shown. For nearly an order of magnitude increase in luminosity, the rate of correlation between trigger towers of significant energy and high-p_T tracks increases by less than 10% in each case. This suggests that track-calorimeter matching will continue to be a powerful tool for background rejection at the highest luminosities.

Table 26 Trigger-tower-track occupancy for 2 GeV and 20 GeV QCD jet k_T and different tower E_T thresholds for low ($4x10^{31}$ $cm^{-2}s^{-1}$) and high luminosity conditions ($5x10^{32}$ $cm^{-2}s^{-1}$).

EM E_T (GeV)	Jet k_T > 2 GeV $4x10^{31}$ $cm^{-2}s^{-1}$	Jet k_T > 20 GeV $4x10^{31}$ $cm^{-2}s^{-1}$	Jet k_T > 2 GeV $5x10^{32}$ $cm^{-2}s^{-1}$	Jet k_T > 20 GeV $5x10^{32}$ $cm^{-2}s^{-1}$
>0.5	9k/197k (4.6%)	42k/161k (26%)	200k/1520k (13%)	92k/291k (33%)
>2	69/297 (23%)	4k/7506 (53%)	1100/3711 (30%)	2130/3482 (61%)
>5	5/9 (50%)	920/1587 (58%)	52/132 (39%)	480/703 (68%)
>10	--	157/273 (58%)	--	96/125 (77%)

The huge numbers of real (and fake) low-momentum tracks in minimum bias events make it impractical to use a track P_T threshold of only 1.5 GeV/c for electron identification. More reasonable values will be found in the range 3-10 GeV/c. Since the rate of fake tracks at these higher momentum thresholds also increases with luminosity, the rate of correlation as a function of track P_T must also be considered.

Table 27 shows such a study, where the fraction of EM object-L1CTT track matches is given as a function of L1CTT track P_T for low k_T jet events at a luminosity of $5x10^{32}$ $cm^{-2}s^{-1}$. These results show that additional rejection is possible by increasing the track P_T and by requiring that the EM object E_T and track P_T match.

Table 27 Trigger tower-track occupancy for a sample of jets with p_T > 2 GeV at 5x10^{32} cm^{-2}s^{-1}. The rate at which tracks of varying P_T are matched to calorimeter trigger towers of increasing E_T thresholds is shown. The entries in the Table are the same as in Table 26.

EM E_T (GeV)	Track P_T >1.5GeV	Track P_T >3GeV	Track P_T >5GeV	Track P_T >10GeV
>0.5	200k/1520k (13.2%)	70k/1520k (4.6%)	30k/1520k (2%)	10k/1520k (0.7%)
>2	1100/3711 (30%)	600/3711 (16.2%)	211/3711 (6%)	60/3711 (2%)
>5	52/132 (39%)	34/132 (26%)	19/132 (14%)	11/132 (8%)
>10	4/12 (30%)	4/12 (30%)	2/12 (20%)	2/12 (20%)

The above studies clearly demonstrate a potential reduction in the EM trigger rate by exploiting the correlations in ϕ and P_T/E_T between L1CTT tracks and calorimeter objects. As mentioned above, the ϕ granularity of EM objects will improve in Run IIb by a factor of 8 (90º quadrants versus 11.25º towers). The increased rejection of the improved ϕ granularity is estimated in Table 28. The EM object – track match fraction is given for two track P_T thresholds and compares quadrant and trigger tower matching. Low k_T jet events at a luminosity of 5x10^{32} cm^{-2}s^{-1} are used as the background sample. We use this study to estimate the increase in background rejection for EM triggers at high luminosity using the L1CalTrack trigger to be an important factor of 2-3.

Table 28 Trigger-tower-track occupancy for a sample of jets with p_T > 2 GeV at 5x10^{32} cm^{-2}s^{-1}. The table presents a comparison of the rate at which tracks of p_T > 1.5 GeV or p_T > 10 GeV are matched to an individual trigger tower or a calorimeter quadrant containing an EM tower above a given threshold. Each line in the table contains the number of matches divided by the total number of quadrants or towers above that E_T threshold.

EM E_T	Track p_T >1.5GeV (quadrants)	p_T> 1.5GeV (towers)	Track p_T > 10GeV (quadrants)	p_T > 10GeV (towers)
2 GeV	2470/3711	1100/3711	225/3711	60/3711
5 GeV	103/132	52/132	21/132	11/132
10 GeV	8/12	4/12	2/12	2/12

5.2.2 Improving L1CTT Rates Using Calorimeter Jets

The previous section presented evidence that the addition of track information can improve the rejection of an electron trigger by requiring a track

close in ϕ to the EM trigger tower. In this section we explore the equally useful converse, namely that the calorimeter can be used to improve the selectivity and background-rejection of tracking triggers. Isolated high P_T tracks are signatures of many types of interesting events. However, the triggers that select these tracks suffer from a large background of fakes, even for a track P_T >10 GeV/c. As has been indicated elsewhere in this document, this problem worsens substantially as the number of multiple interactions increases. The matching of these tracks to signals in the calorimeter has the ability to confirm the existence of the tracks themselves, and also to verify their momentum measurement.

In this study, our matching algorithm considers individual L1CFT sectors ($\Delta\phi$=4.5°) with at least one track of a given minimum P_T, and matches them in ϕ to whatever trigger towers they overlap. By doing this, we avoid double counting some of the redundant track solutions that cluster near to each other. In about one third of the sectors, these tracks will overlap two different trigger towers in ϕ; each match is counted separately. The results of track-trigger tower matching are given in Table 29 using the low k_T, high luminosity QCD sample as representative background. Note that for this study, the E_T in the table is the Total E_T (EM+EH), not just the EM E_T. Given that most tracks are hadrons, this is more representative of the true energy that should be matched to a given track.

Table 30 Trigger-tower-track matching for a sample of jets with k_T > 2 GeV at $5x10^{32}$ $cm^{-2}s^{-1}$. The number of CFT trigger sectors containing at least one track above a given p_T threshold is shown, both without and with matching to calorimeter trigger towers of increasing total E_T.

track p_T	# sectors with tracks	Tot E_T > 1 GeV	> 2 GeV	> 5 GeV	> 10 GeV
> 1.5 GeV	52991	16252	3218	200	13
> 3 GeV	12818	5188	1529	144	13
> 5 GeV	4705	1562	476	73	9
> 10 GeV	2243	655	141	31	5

With this algorithm we find substantial rejections from even mild trigger tower thresholds. For example, a 10 GeV/c track matching to a 5 GeV trigger tower provides a factor of ~70 rejection against fakes. Matching any track to a 2 GeV tower provides approximately a factor of 10 rejection. The rejection shown in this table is essentially sufficient to allow the high-p_T single and di-track triggers to function at the highest luminosities. At this preliminary stage this is a very promising result.

Clearly further simulation work is needed and is in progress. Additionally we must include the CPS and FPS elements, which should provide additional rejection. Finally, simulation results must be cross-checked with results from

collider data. These latter studies will be carried out in the near future when the L1CTT trigger is operational.

5.3 Implementation

Our working assumption is that the L1CalTrack trigger architecture can be made identical to that of the existing and operational Level 1 Muon (L1Muo) trigger with only small and straightforward modifications. Specifically, we will use minimally modified muon trigger (MTCxx) cards with a new flavor board (MTFB) that performs the calorimeter-track match algorithms. However even the new flavor board will be a straightforward upgrade of the existing flavor board that contains the muon detector-track match algorithms. The L1CalTrack trigger crate manager (MTCM) and trigger manager (MTM) will be duplicates of those used for the L1Muo trigger. Most importantly, the engineering effort on traditionally time-consuming details such as synchronization, buffering, messages to L2 and L3, electronics testing, monitoring, power supplies, and crate infrastructure is then virtually nil. A key (and unresolved) question is whether 16 serial link inputs will suffice for the L1CalTrack trigger. Nevertheless, given all the advantages of using L1Muo trigger hardware and the fact that the L1Muo trigger is being successfully operated, we continue on this path at present.

5.3.1 L1Muo System

A brief description of the L1Muo trigger is given here. This is followed by a few technical details on serial links, synchronization, and buffering. A brief description of how the L1CalTrack trigger uses the L1Muo hardware and possible need modifications.

The L1Muo trigger satisfies the following requirements:

- Delivers an L1Muo Trigger decision to the TF at 3.3 μs after Bunch Crossing (BC)
- Transmits an L1Muo decision for every BC not occurring in the Synch Gap
- Operates with 132 or 396 ns BC times
- Synchronizes inputs to each Muon Trigger Card (MTCxx)
- Provides buffering for input and output data pending an L1 decision from the TF
- Provides 16 buffers for data pending a Level 2 (L2) decision from the TF
- Provides 8 buffers for data pending readout to Level 3 (L3)
- Deadtimeless operation
- Field programmable trigger logic
- Online and offline monitoring
- Complete documentation

Thus the L1CalTrack trigger satisfies the same requirements.

A block diagram of the L1Muo Trigger is shown Figure 65. There are three custom VME crates of Muon Trigger Cards (MTCxx's) corresponding to the central (CF), north (EFN), and south (EFS) geographic regions of the DØ detector. There is one custom VME crate that serves as a Muon Trigger Manager (MTM). The VME crates reside on the detector platform and are thus inaccessible during data-taking.

For each crossing, data is transmitted from muon detector front-end cards and the L1CTT trigger to the Muon Trigger Cards (MTCxx's). The information is transmitted at 1060 Mbits/s over coaxial cable using the AMCC S2032/2033 serial link chip set implemented on Serial Link Daughter Boards (SLDB's). Within each geographic region, the MTCxx cards form local trigger decisions for each octant. The actual trigger decision logic is implemented in Altera ACEX series FPGA's contained on a Muon Trigger Flavor Board (MTFB) that plugs into each MTCxx. Currently we have two types of MTFB's called 05 and 10. The first matches tracks from the L1CTT with hits in the muon detector scintillation counters while the second finds tracks using the muon detector wire chambers. A photo of the MTCxx card is shown in Figure 66.

Figure 66The octant trigger decisions are sent over a custom VME backplane to the Muon Trigger Crate Manager (MTCM) which subsequently forms trigger decisions for each geographic region. The regional trigger decisions are then transmitted by the MTCM's using Gbit/s serial links to the Muon Trigger Manager (MTM) which forms the global muon trigger decision that is sent to the Trigger Framework (TF). There are 256 L1Muo trigger terms that the user can choose from at begin run time. L1MUO triggers can be chosen based on geographic region (including $1.5 < |\square| < 2.0$ where there is no L1CFT Trigger coverage), multiplicity (0-3), PT threshold (presently 2, 4, 7, and 11 GeV/c), and quality (called Loose and Tight). Presently there are 32 user defined AND-OR terms sent to the TF representing the global L1Muo trigger decision. A photo of the MTCM card is shown in Figure 67.

On receipt of an L1 Accept from the TF, the L1Muo Trigger sends its trigger decision and additional information to the L2 trigger system (via the Serial Link Interface Cards (SLIC's). On receipt of an L2 Accept, the L1Muo trigger sends its trigger decision and additional information to the L3 trigger system (via the Muon Readout Cards (MRC's). Additionally, the L1Muo trigger may send all of its input data to the L3 trigger system for 1 of N beam crossings (where N is user-defined).

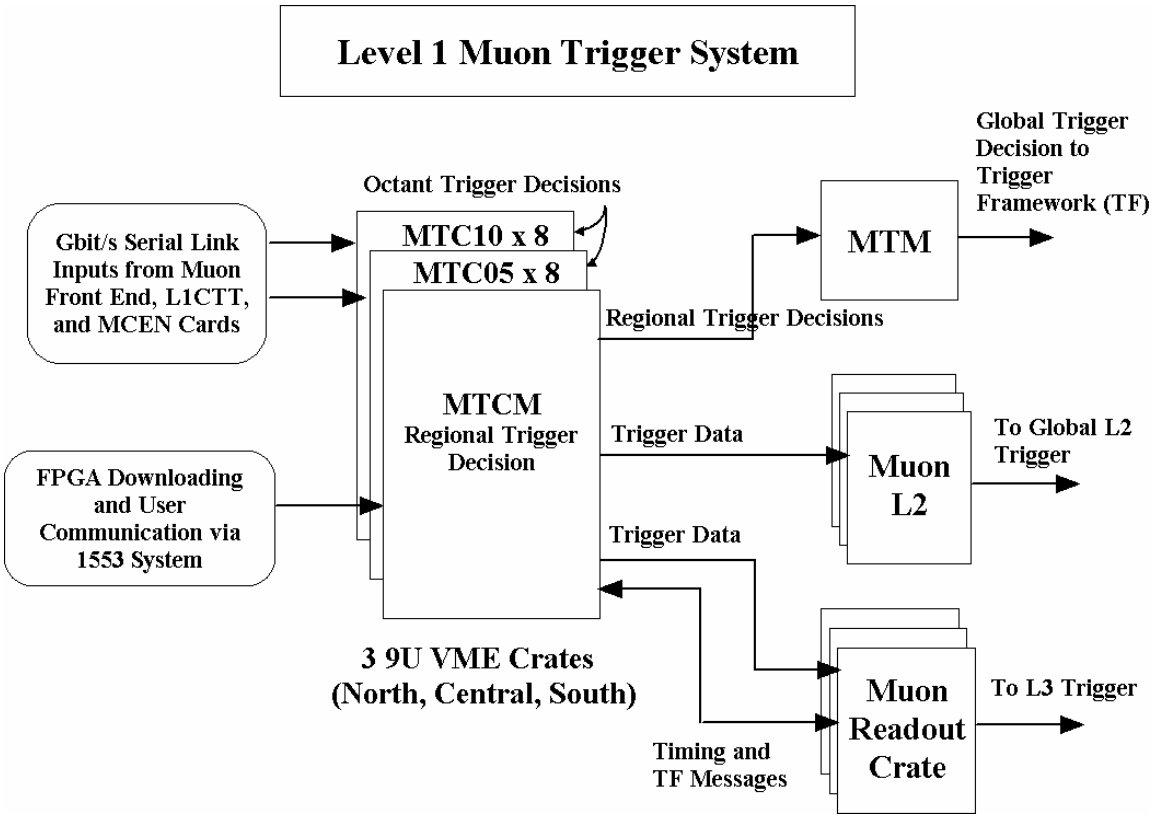

Figure 65. Block diagram of the L1Muo trigger system.

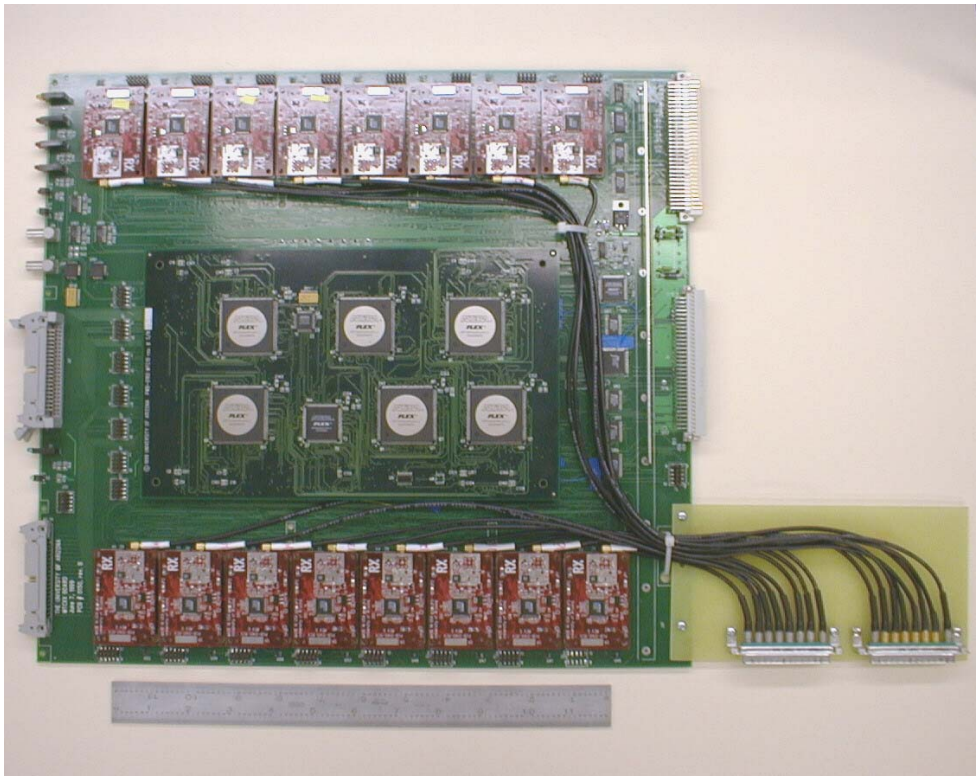

Figure 66. Photo of the MTCxx card for the L1Muo trigger system.

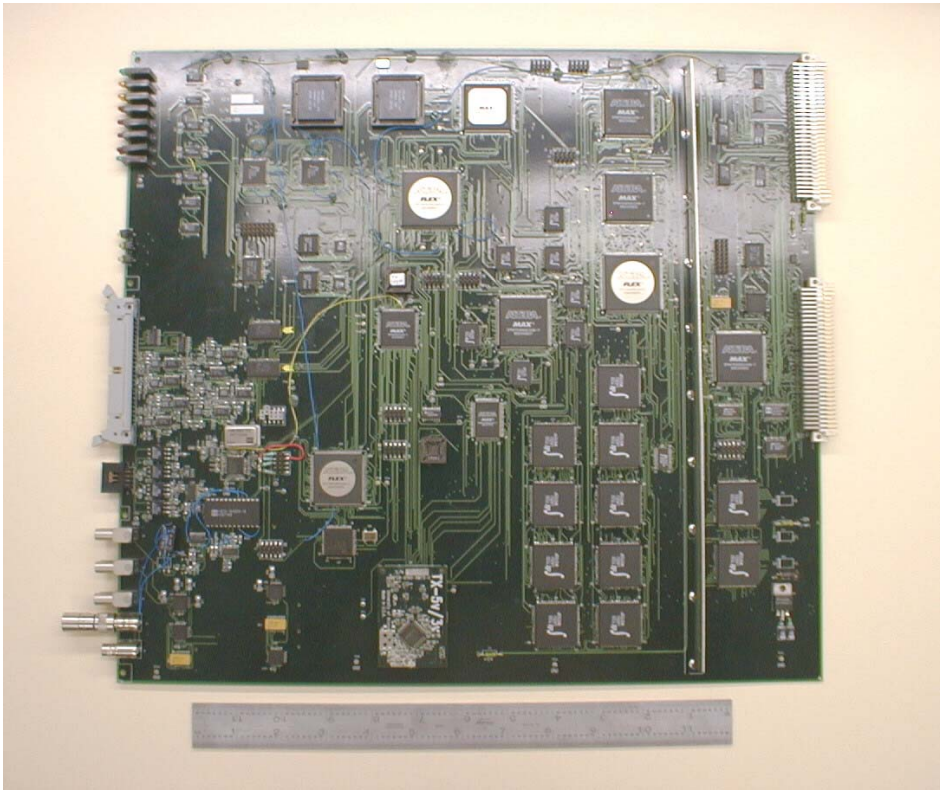

Figure 67. Photo of the MTCM card for the L1Muo trigger system.

Timing and trigger information from the Trigger Framework (TF) takes the following path. The TF sends this information to the Muon Fanout Card (MFC) via a Serial Command Link (SCL). The MFC distributes this information to MRC cards over the VME backplane. An MRC subsequently sends this information to each Muon Trigger Crate Manager (MTCM). The MTCM distributes this information to the MTCxx cards over the VME backplane. Users may communicate with the L1Muo Trigger system via two paths: one is the MIL-STD-1553B system and the other a UART between MTCM and MRC.

5.3.1.1 Serial Links

One of the key elements of the L1Muo Trigger system are the Serial Link Daughter Boards (SLDB)'s. Gbit/s serial transmission over coaxial cable was chosen to maximize the amount of information that could be brought onto each MTCxx card, to minimize the cable plant, and to minimize the cost. At the time of the MTCxx design and even today, transmission over optical fiber would have been prohibitively expensive given our need for 768 serial links for the L1MU Trigger and 768 serial links for the MCEN system.

The chipset chosen is the AMCC 2042/2043 which is Fiber Channel compatible. The SLDB also contains an Altera 7K series EPLD which handles 8b/10b encoding and parity calculation on the transmitter and 8b/10b decoding and parity checking on the receiver. The MTFB receiver also contains an equalizer circuit and amplifier (HP IVA-05208) needed for error-free transmission over ~150 feet of coaxial cable (LMR-200). Block diagrams of the SLDB transmitter and receiver are shown in Figure 68 and Figure 69. Eye patterns before and after equalization are shown in Figure 70and Figure 71.

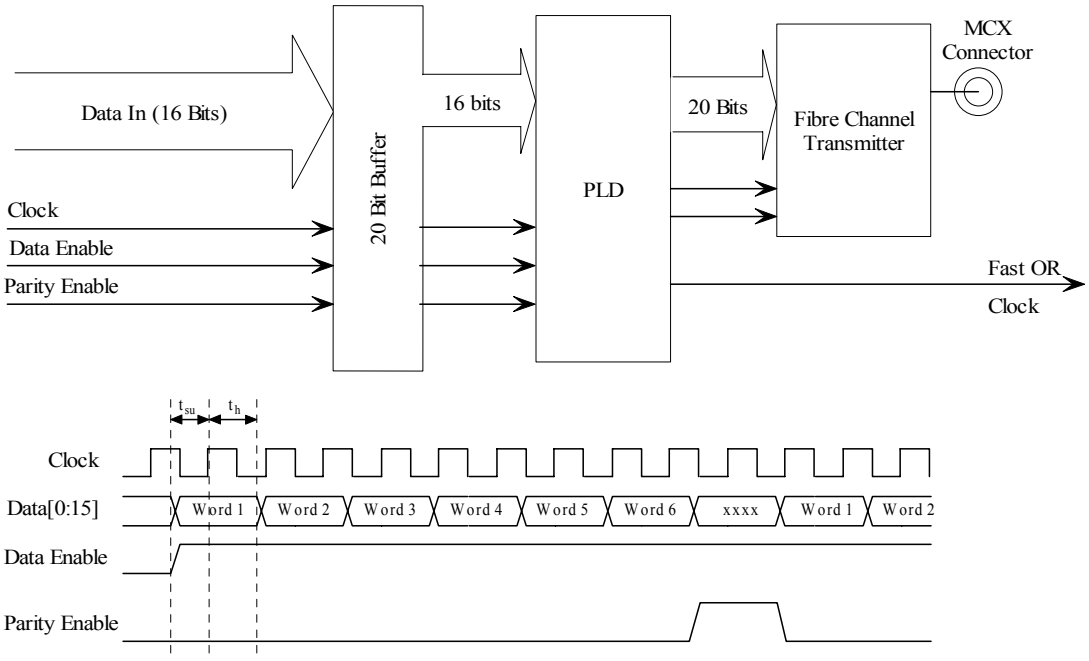

Figure 68. Block diagram of the SLDB transmitter.

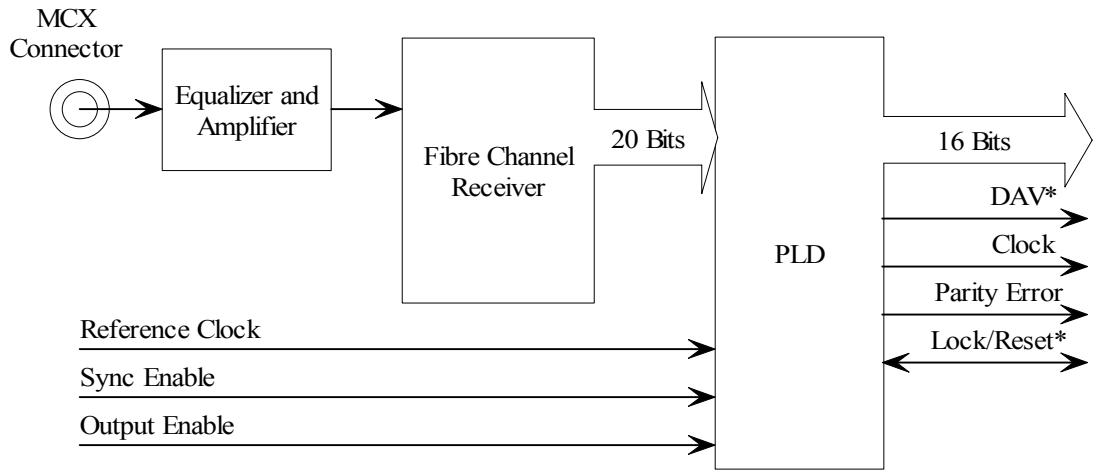

Figure 69. Block diagram of the SLDB receiver.

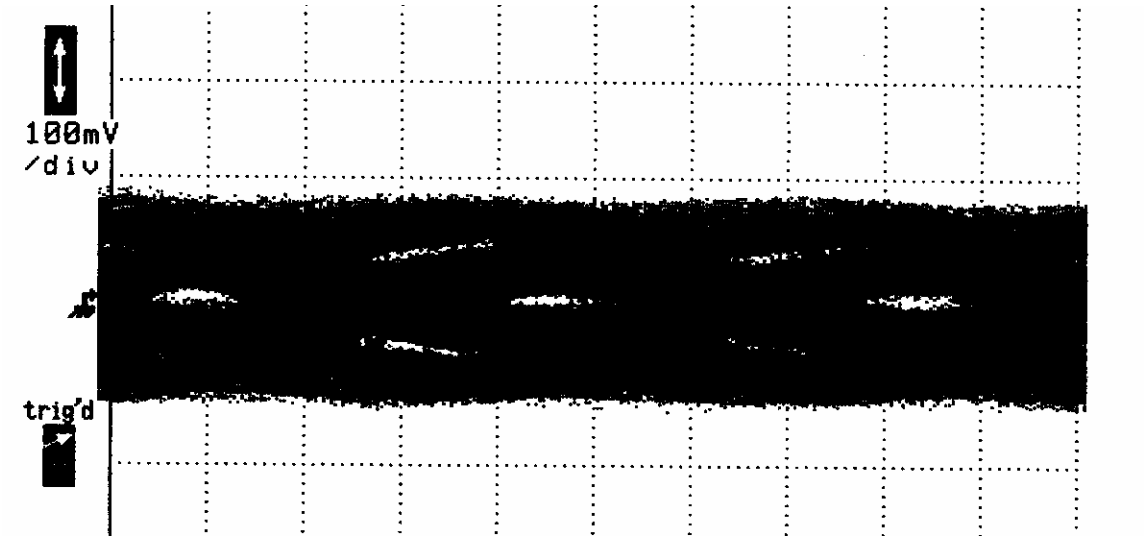

Figure 70. Eye pattern for transmission over 150 feet of LMR-200 coaxial cable without equalization and amplification.

Figure 71. Eye pattern for transmission over 150 feet of LMR-200 coaxial cable with equalization and amplification.

The SLDB's have been running in the L1Muo trigger system at DØ for well over a year and have experienced no problems. Additionally the TF uses these SLDB's to transmit SCL information to all Geographic Sectors. Again, their operation has been problem free.

5.3.1.2 Synchronization

A common question is how is the data from different input cables to the MTCxx cards synchronized. The answer is by the use of FIFO's. After INIT, all input FIFO's (Figure 72) are reset. In addition, all active SLDB receivers should be receiving K28.5 Idle characters that are not written to the FIFO's. Thus all FIFO's remain empty.

After INIT, data transmission begins. As data is received at each of the sixteen SLDB receivers, it is immediately written to the input FIFO's. When all FIFO's have data, the MTFB is told to begin processing. Also, an Input Ready signal is sent to the MTCM alerting it that trigger processing has begun. So when all of the Input FIFO's have gone non-empty it is guaranteed that all the input data from the first beam crossing has been synchronized and assembled.

The trigger decision output of the MTFB's are written to FIFO's as well. Each MTCxx sends a Data Ready to the MTCM alerting it that trigger processing is complete. When all MTCxx's have asserted Data Ready, the MTCM returns Start Processing to the MTCxx's which send their MTFB trigger decision data to the MTCM.

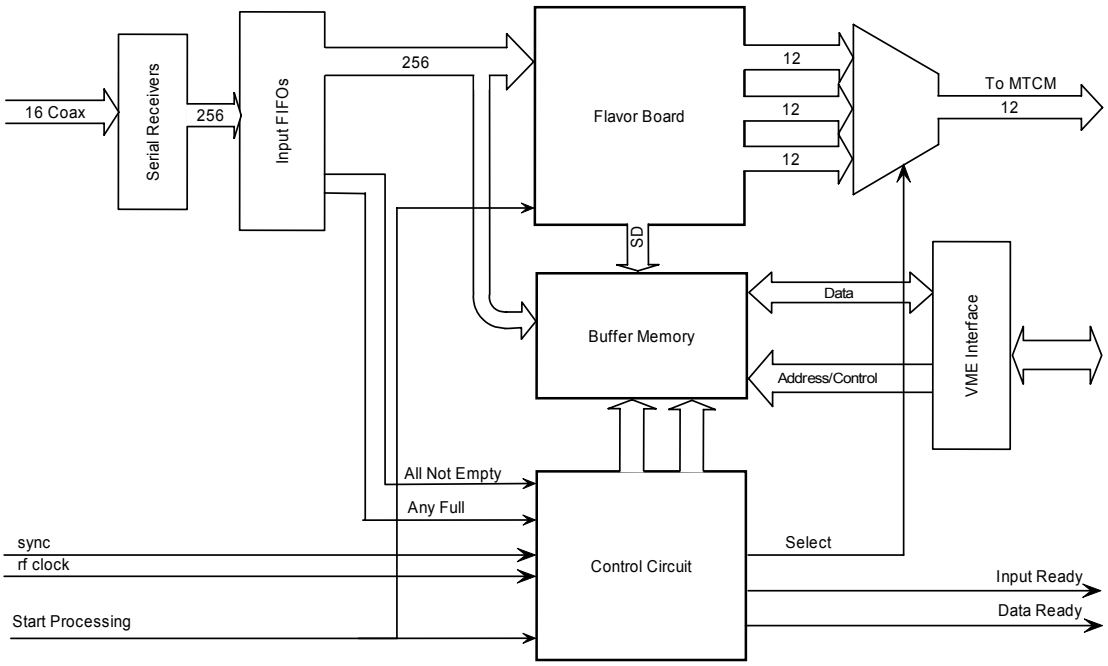

Figure 72. Block diagram of the MTCxx trigger processing.

5.3.1.3 Buffering

Buffering of data pending L1 trigger, L2 trigger, and L2 transfer decisions is achieved through the use of dual port memory (DPM) and a set of FIFO's that keep track of the DPM addresses (see Figure 73). Input data, octant trigger decision data, regional trigger decision data, and other (status, error) information is the data that must be buffered. After Initialization, the appropriate DPM addresses are listed in an Empty Buffer FIFO. With the arrival of the first event, input data, trigger decision data, and other information are written to DPM. The starting address of this DPM data is written to the L1 Pending Buffer FIFO. On receipt of an L1 Accept, the DPM address is written to the L2 Pending Buffer FIFO. A mirror of this FIFO is used by the MTCM to read the DPM in forming the L2 Data message. For L1 Rejects, the DPM address is returned to the Empty Buffer FIFO. On receipt of an L2 Accept, the DPM address is written to the Buffer Transfer FIFO. This is the address used by the MTCM to read the DPM in forming the L3 Data message. For L2 rejects, the DPM address is returned to the Empty Buffer FIFO.

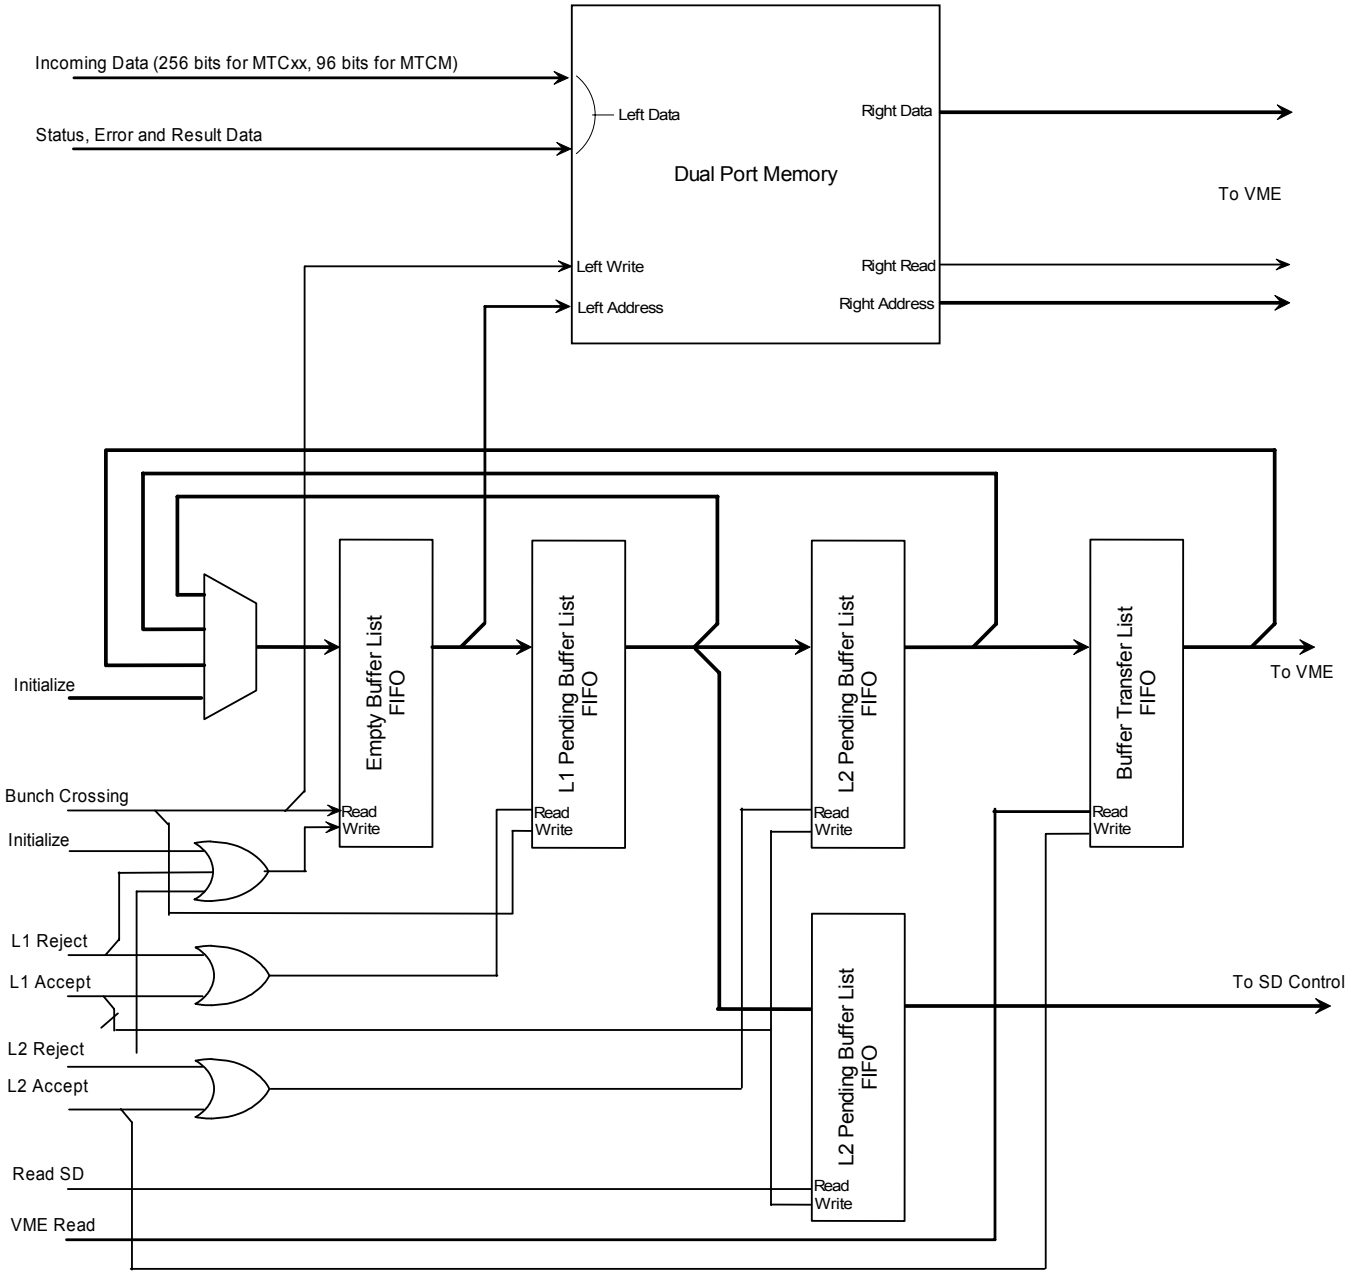

Figure 73. Block diagram of the L1Muo buffering scheme.

5.4 L1CalTrack Trigger

A block diagram of the L1CalTrk trigger is shown in Figure 74. Note the many similarities to Figure 65. Differences between the L1CalTrk and L1Muo triggers include details and source of the inputs, the MTFB and algorithms that implement the calorimeter-track matching, timing, and the location of the L1CalTrk crates.

Each MTCxx card can accept up to sixteen serial inputs. A preliminary definition of these inputs for the L1CalTrack trigger is given in Table 31. For

each crossing, each input effectively contains up to 96 bits (six 16 bit words at 53 MHz) of information. The bit assignment for the L1CTT tracks is defined. The bit assignments for the L1Cal and L1FPS triggers are not yet defined. Note that if additional inputs are needed on the MTCxx card, two to four additional inputs could be accommodated. This would involve changes to the existing MTCxx design and layout however. A requirement of the Run IIb L1CTT trigger is to include the CPS information in the tracks sent to L1Muo. This is because the tracks sent to the L1CalTrack trigger are simply sent on the one of the dual outputs of the L1CTT's SDLB's. Preliminary discussions with the relevant engineers shows including this information to be straightforward.

Table 31. Cable inputs for MTCcal cards.

Cable #	Bits / Cable	Definition	Source
1-10	6 tracks x 16 bits / track	L1CTT trigger tracks (includes CPS)	L1CTT
11-12	96	EM and Jet objects	L1Cal
13-16	96	FPS shower objects	L1FPS

The octant decision formed by MTCcal cards consists of 36 bits that are subsequently sent over the backplane to the MTCM card. The definition of these bits is presently undefined. Note that 36 bits is a hard limit so we must ensure this is sufficient for the variety of triggers produced by the MTCcal cards.

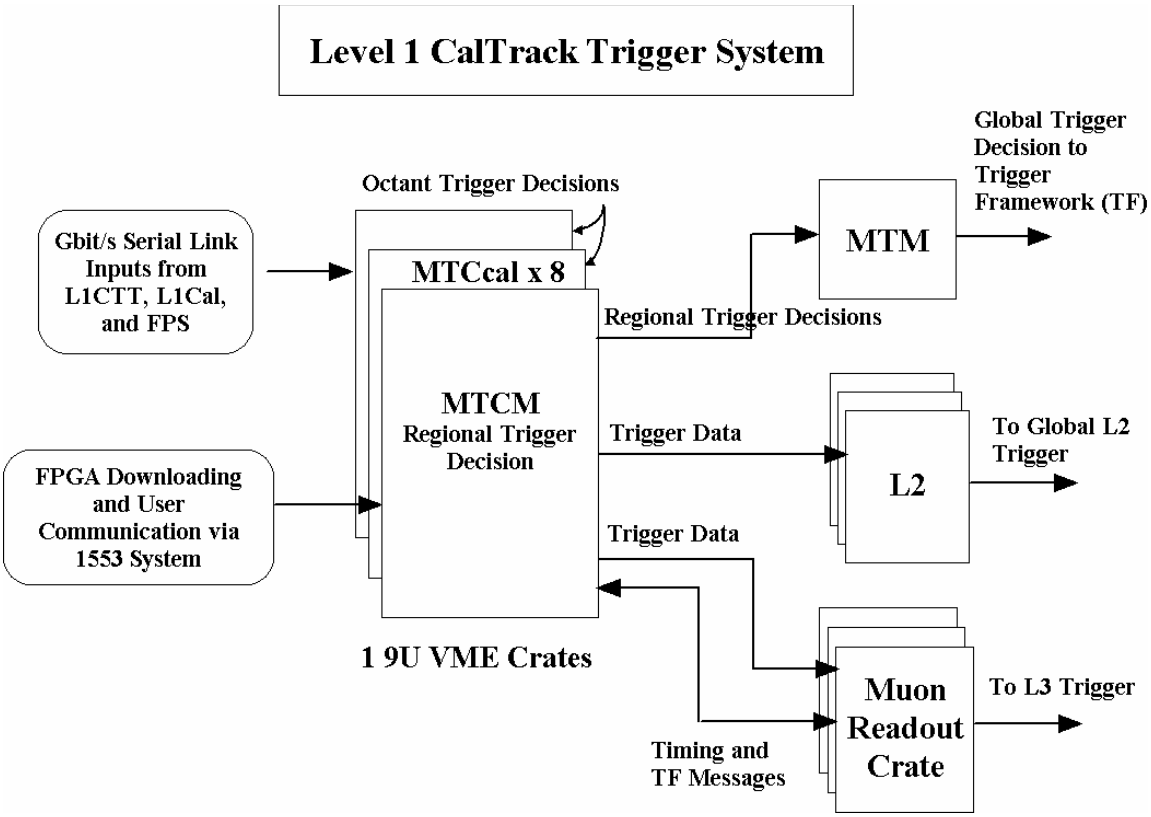

Figure 74. Block diagram of the L1CalTrack system.

Each MTCM receives timing and trigger information from the trigger framework (TF) via a Muon Readout Card (MRC). The detailed path is from TF via a Serial Command Link (SCL) to an Muon Fanout Card (MFC) that distributes the information over a custom VME backplane to the MRC's. The connections between MRC and MTCM are given in Table 32 and Table 33. Because these signals will not be transported over a standard 24-wide Astro cable, a specialized cable between the MTCM and MRC will be needed.

Table 32. Timing and data signal definitions between the MTCM and MRC. The signals are sent over coaxial "Astro" cable.

Pin	Definition	Pin	Definition
1	L3 Data +	2	L3 Data -
3	RF Clock +	4	RF Clock -
5	Encoded Timing +	6	Encoded Timing -
7	L2 Data +	8	L2 Data -

Table 33. Trigger information definitions between the MTCM and MRC. These signals are transmitted over 50c twist and flat cable.

Pin	Definition	Pin	Definition
1	BC Number 1 +	2	BC Number 1 -
3	BC Number 2 +	4	BC Number 2 -
5	BC Number 3 +	6	BC Number 3 -
7	BC Number 4 +	8	BC Number 4 -
9	BC Number 5 +	10	BC Number 5 -
11	BC Number 6 +	12	BC Number 6 -
13	BC Number 7 +	14	BC Number 7 -
15	BC Number 8 +	16	BC Number 8 -
17	INIT +	18	INIT -
19	L1 Accept +	20	L1 Accept -
21	L2 Error +	22	L2 Error -
23	L2 Accept +	24	L2 Accept -
25	L2 Reject +	26	L2 Reject -
27	UART Transmit +	28	UART Transmit -
29	Buffer Available +	30	Buffer Available -
31	Strobe +	32	Strobe -
33	UART Receive +	34	UART Receive -
35	L1 Error +	36	L1 Error -
37	L1 Busy +	38	L1 Busy -
39	L2 Busy +	40	L2 Busy -
41	GND	42	GND
43	GND	44	GND
45	GND	46	GND
47	GND	48	GND
49	GND	50	GND

Each MTCM is connected to the Muon Trigger Manager (MTM) via a serial link. As an aside, the MTM is made of an MTCxx card with an MTM MTFB. The format of this (presently) bitwise transfer is given in Table 34.

Table 34. MTCM serial link data sent to the MTM.

Word #	Definition
0	Spare
1	MTCM trigger decision bits 0-11
2	MTCM trigger decision bits 12-23
3	MTCM trigger decision bits 24-35
4	Spare
5	Parity

On receipt of an L1 Accept, the MTCM transmits data found in Table 35 to the L2 Muon Trigger. The data block size is the same for all events. Presently sixteen-bit words are sent with the lower byte first. Finally note there is a similar but not identical block of data from the MTM crate.

Table 35 MTCM output to L2.

Byte #	Definition
	BeginTransmission (K28.0)
1-2	Word Count (=35)
3-4	Module ID
5-6	Local Crossing Number
7-8	Local Turn Number
9-10	Event Status (undefined)
11-12	Event Status (undefined)
13-14	MTCM Status/Control Register
15-16	MTCM Event Error Register
17-22	MTCM Regional Trigger Decision
23-28	MTCcal Octant 0 Trigger Decision
29-34	MTCcal Octant 1 Trigger Decision
35-40	MTCcal Octant 2 Trigger Decision
41-46	MTCcal Octant 3 Trigger Decision
47-52	MTCcal Octant 4 Trigger Decision
53-58	MTCcal Octant 5 Trigger Decision
59-64	MTCcal Octant 6 Trigger Decision
65-70	MTCcal Octant 7 Trigger Decision
	End Transmission (K23.7)

On receipt of an L2 Accept, the MTCM transmits data found in Table 36. to the MRC where it is subsequently read by the VBD. There are two possible data block sizes. Recall the L1Muo trigger system has the option of transmitting all of its input data for 1 of N events (where N is user defined). The two block sizes correspond to whether or not the input data block is included. Presently sixteen-bit words are sent with the lower byte first. Finally note there is a similar but not identical block of data from the MTM crate.

Table 36. MTCM output sent to L3.

Word #	Definition
	BeginTransmission (K28.0)
	Upper 16 bits of VBD Word Count = 0
	VBD Word Count = 246 or 1142 decimal
1	Word Count = 492 or 2284 decimal
2	Module ID
3	Local Crossing Number
4	Local Turn Number
5	Event Status = MTCM Event Error Register from DPM
6	Event Status (undefined)
7	MTCM Control Register
8	MTCM Event Error Register from DPM
9	MTCM Latched Error Register
10	MTCM Trigger Logic FPGA Program ID
11	MTCM Message Builder FPGA Program ID
12	MTCM L1_Accept Divider Number
13	MTCM L2_Accept Divider Number
14	MTCM MTCxx Readout Mask
15	MTCM MTCxx Trigger Mask
16	MTCM L1 Error Mask
17-19	MTCM Regional Trigger Decision
20	MTCxx #0 MTCxx Serial Number
21	MTCxx #0 Flavor Board Type/Serial Number
22	MTCxx #0 MTCxx Status Register
23	MTCxx #0 MTCxx Flash Memory Status
24-34	MTCxx #0 Various Error Registers (82,86,8a,8e,92,94, pad rest with zeros)
35-38	MTCxx #0 Various Mask Registers (06,08,0a,0c)
39-46	MTCxx #0 FPGA Program ID's (total of 8)
47-73	MTCxx #1 Info Block
74-100	MTCxx #2 Info Block
101-127	MTCxx #3 Info Block
128-154	MTCxx #4 Info Block
155-181	MTCxx #5 Info Block
182-208	MTCxx #6 Info Block

209-235	MTCxx #7 Info Block
	MTCcal Octant 0 Trigger Decision
	MTCcal Octant 1 Trigger Decision
	MTCcal Octant 2 Trigger Decision
	MTCcal Octant 3 Trigger Decision
	MTCcal Octant 4 Trigger Decision
	MTCcal Octant 5 Trigger Decision
	MTCcal Octant 6 Trigger Decision
	MTCcal Octant 7 Trigger Decision
	MTCxx Input Data
	End Transmission (K23.7)

Finally, the MTM MTFB (which is used in concert with an MTCxx mothercard) takes the regional trigger decision data from the three MTCM cards and forms a variety of global trigger decisions. The specific triggers sent to the Trigger Framework are downloaded during the physics trigger download of a data acquisition run. The data sent to the Trigger Framework from the MTM MTFB is given in Table 37.

Table 37. MTM data sent to the Trigger Framework.

Pin	Definition	Pin	Definition
1	Specific Trigger 0 +	2	Specific Trigger 0 -
3	Specific Trigger 1 +	4	Specific Trigger 1 -
5	Specific Trigger 2 +	6	Specific Trigger 2 -
7	Specific Trigger 3 +	8	Specific Trigger 3 -
9	Specific Trigger 4 +	10	Specific Trigger 4 -
11	Specific Trigger 5 +	12	Specific Trigger 5 -
13	Specific Trigger 6 +	14	Specific Trigger 6 -
15	Specific Trigger 7 +	16	Specific Trigger 7 -
17	Specific Trigger 8 +	18	Specific Trigger 8 -
19	Specific Trigger 9 +	20	Specific Trigger 9 -
21	Specific Trigger 10 +	22	Specific Trigger 10 -
23	Specific Trigger 11 +	24	Specific Trigger 11 -
25	Specific Trigger 12 +	26	Specific Trigger 12 -
27	Specific Trigger 13 +	28	Specific Trigger 13 -
29	Specific Trigger 14 +	30	Specific Trigger 14 -
31	Specific Trigger 15+	32	Specific Trigger 15 -
33	Gap +	34	Gap -
35	Ground	36	Ground
37	Strobe +	38	Strobe -
39	Ground	40	Ground

5.5 L1CalTrack Cards

In this section we briefly summarize the functions of VME cards and daughter boards comprising the L1CalTrk Trigger.

5.5.1 Serial Link Daughter Board (SLDB)

The Serial Link Daughter Boards (SLDB's) are used to transmit data over coaxial cable between the L1CTT, L1Cal, and L1FPS triggers (transmitters) and the MTCxx cards (receivers) in the L1CalTrack trigger. There are typically one or two SLDB's on the L1CTT, L1Cal, and L1FPS triggers and sixteen SLDB's on the MTCxx cards. The serial links used are the Fiber Channel compatible AMCC S2042/S2043 chipsets. Altera 7K series EPLD's are used for 8b-10b encoding/decoding and parity checking. Seven sixteen-bit words are transmitted at 53 MHz. With 8b-10b encoding this gives a serial transfer rate of 1060 Mbits/s. Error free data transmission over 150 foot lengths of coaxial cable

(LMR-200) is achieved by using an equalization circuit and high-speed amplifier (HP IVA-05208) on the SLDB receiver boards.

5.5.2 Muon Trigger Card (MTCxx)

The Muon Trigger Cards (MTCxx cards) are the primary trigger cards in the L1CalTrack system. The MTCxx card is a generic VME card that accepts a Muon Trigger Flavor Board (MTFB) as a daughter card. The actual calorimeter-track match trigger logic is implemented on the MTFB. The MTCxx/MTFB combination is used form an octant trigger decision. The MTCxx card also contains sixteen SLDB receivers that are used to accept data from the various input sources. The primary functions of the MTCxx cards are to receive and synchronize sixteen serial inputs, to buffer input data (which is subsequently sent to the MTFB) and to buffer input and supplemental trigger decision data pending L1, L2 and L3 accepts.

5.5.3 Muon Trigger Flavor Board (MTFB)

The Muon Trigger Flavor Board (MTFB) is a daughterboard that is used in concert with the MTCxx card. For the L1CalTrack trigger a new flavor board will be used called MTCcal. This flavor board will contain the calorimeter-track match trigger logic. Note that the present MTC05 MTFB matches tracks from the L1CTT with hits in the muon detector scintillation counters and hence can be used as a good estimate for the calorimeter-track logic. The MTCcal MTFB will match tracks from the L1CTT with EM and jet objects from L1Cal. While we must simulate all calorimeter-track match algorithms in both the C++ trigger simulator and MAXPLUS2 FPGA simulator before a final choice of FPGA family is made, preliminary studies show that the four FPGA's on the present MTC05 MTFB can be replaced with one larger FPGA. In addition, the MTCcal MTFB will contain a Gbit/s serial link. This will allow the MTCM to be bypassed in forming the trigger decision that is sent to the Trigger Framework.

5.5.4 Muon Trigger Crate Manager (MTCM)

The Muon Trigger Crate Manager (MTCM) reads the octant trigger decisions from each MTCxx card and uses this data to form a regional trigger decision that is subsequently sent to the Muon Trigger Manager (MTM). The MTCM also serves to buffer the octant trigger decision data from each MTCxx card and the regional trigger decision pending L2 and L3 accepts. L2 data and L3 data are sent to their respective systems using the Cypress Hotlink (CY7B923/CY7B933) chipset. The MTCM accepts timing and trigger information from the Trigger Framework (TF) and reports error and busy conditions to the TF also via the MRC. The MTCM maintains two paths by which a user can communicate to the L1MU system: 1553 and UART. The UART path is a link between the MTCM and MRC. Either path can be used for downloading, monitoring, and testing.

5.5.5 Muon Splitter Cards (MSPLIT)

The MSPLIT cards are used as fanouts for many of the Gbit/s serial input signals to the MTCxx cards. They are mentioned here because the location of the L1CalTrack crates is now in MCH1 rather than collision hall. This means we

must send the SLDB Gbit/s data over 180-200 feet of coaxial cable. This is at the upper limit of our demonstrated length of guaranteed error free transmission. Should an additional signal boost be required, we can use the MSPLIT cards just inside of MCH1. The splitters are implemented on 9U VME cards but without any VME functionality. They consist of a high-speed amplifier (HP IVA-05208) and high bandwidth transformer splitter. Each card has 8 inputs and up to 24 outputs. Hence they can be used as cheap repeaters.

5.5.6 Muon Trigger Test Card (MTT)

The purpose of the Muon Trigger Test (MTT) card is to test and debug the L1CalTrack trigger cards. These include the MTCxx and MTCM cards. The MTT simulates front-end data for the MTCxx cards using sixteen SLDB transmitters with user-defined input. The MTT serves as an MRC in that it generates trigger and timing data used by the MTCM cards. The MTT also contains three Hotlink receivers so that the L2 and L3 data outputs of the MTCM can be tested. It is mentioned as another example of how engineering is greatly minimized by using existing hardware and/or designs.

5.5.7 Timing

A summary spreadsheet of the total latency of the L1CalTrack trigger is given in Table 38. A detailed spread sheet is also available. The spreadsheet assumes that the Run IIb Level 1 trigger latency will be increased by six BC over that for Run IIa (from 3300 to 4092ns). The pipelines in the existing front ends of the silicon, fiber tracker, and calorimeter can accommodate this additional latency. Preliminary engineering studies show that the pipelines in the PDT and muon scintillator front ends can be increased by making modest changes to the existing electronics. Thus the increased latency can be accommodated by all front end systems. This increase will provide more than 600ns of headroom in the L1CalTrack trigger latency. The spreadsheet also assumes that the MTCM can be bypassed in the trigger decision chain from MTCxx card to MTM. It also assumes a savings of 132 ns in the existing MTC05 logic that has been identified and simulated in MaxPlusII.

Table 38. Trigger latency of the L1CalTrack trigger.

Element	Time (ns)
BC to ADF card (measured)	650
ADF processing and serialization	1147
TAB processing	728
MTCxx processing (measured)	729
Total Latency	3254
TF L1 Decision Time	4092
Difference	-838

6 Level 2 βeta Trigger

6.1 Motivation

An overview of the Level 1 and Level 2 trigger system is given in Section 2.1. At Level 2, preprocessors analyze output from the Level 1 trigger of each detector system in parallel. (The L2 Muon and L2 Silicon Tracking Trigger (L2STT) preprocessors also receive fully digitized detector data.) Event selection takes place based on the high level (physics object) information available in a Global stage after detector preprocessing is complete. The preprocessors instrumented in the Run IIb trigger include calorimeter, preshower, silicon tracker, and muon components.

The input rate to the L2 trigger is limited by the SMT digitization deadtime, and the output rate is limited by the calorimeter precision readout deadtime. Since both limits are constant from Run IIa to Run IIb, the primary charge for Level 2 will be to maintain Run IIa rejection levels (factor of ~5) within the same time budget (to fully realized the advantages from our L1 enhancements). Maintaining Level 2 rejection in the Run IIb trigger will be more challenging as some algorithms used in the Run IIa Level 2 trigger move upstream to Level 1 for Run IIb. To accomplish its goal, Level 2 must make better use of the time budget by using more powerful processors. This project is already under way with the construction of the 'Level 2 βeta' processors, initially conceived to deal with lower than expected production yields in the Run IIa Alpha processors and to offer a clear upgrade path for increases in future performance.

6.2 L2βeta Architecture

All L2 processors occupy 9U VME64 for physics crates. These crates provide dual backplanes: a standard VME bus, and a custom-built 128-bit "Magic Bus" or MBus (a handshaking bus capable of data transfer rates up to 320 MB/s). Each crate contains a number of devices for communication with the experiment's front end and trigger systems and at least two processor cards for analysis of detector subsystem data. The processors are configured for Administrator or Worker functions. Where appropriate, additional specialized hardware for data conversion or processing are included (L2Muon, L2STT). A Worker node applies trigger algorithms to its input data. The Administrator does all event processing and local trigger control tasks that do not involve the application of the trigger algorithm. These include verifying data integrity, controlling communication with the trigger framework, controlling the output of monitoring data, and controlling the readout of events to the higher trigger levels.

The L2β processors[11] rely on commercially produced single board computers (SBCs). Each SBC resides on a 6U CompactPCI (cPCI) card providing access to a 64-bit, 33/66 MHz PCI bus via its rear edge connectors. Such cards are

[11] A detailed L2β TDR is available at http://galileo.phys.virginia.edu/~rjh2j/l2beta/

currently available "off the shelf" from several vendors including Advantech-nc[12], VMIC[13], Diversified Technology Inc.[14], and Teknor[15]. The remaining functionality of the board is implemented in a large FPGA and Universe II[16] VME interface mounted on a 6U-to-9U VME adapter card as shown in Figure 75 - Figure 76.

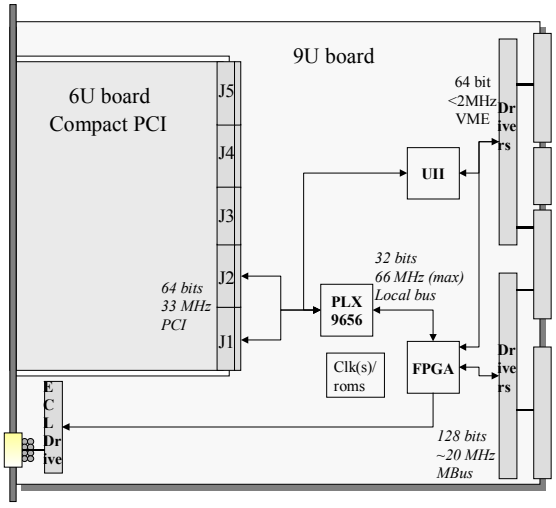

Figure 75. Model for the L2β processor card. Connectors J1 and J2 provide a 64-bit cPCI connection to the CPU.

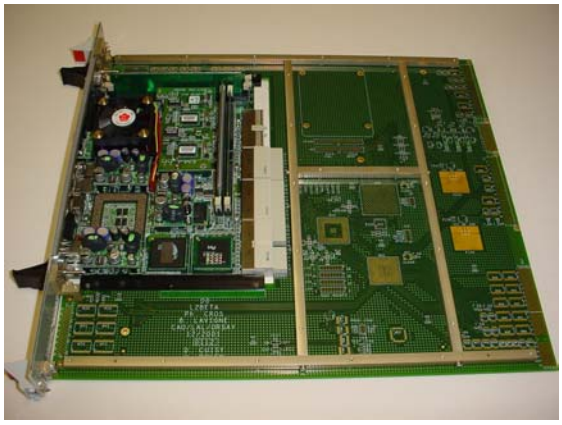

Figure 76. Level 2 βeta prototype. The SBC and 9U cards are shown with the mechanical assembly. The 9U card is shown without components installed.

The adapter card contains all DØ-specific hardware for Magic Bus and trigger framework connections. Custom I/O functions on this card will be implemented in a single FPGA (Xilinx XCV405E) plus assorted logic converters and drivers. This device is particularly suited to our application, because of its large amount of available Block RAM. 70KB of RAM (in addition to >10K logic cells) is used to implement internal data FIFOs and address translation tables for broadcasting data from the Magic bus to CPU memory, reducing the complexity of the 9U PCB. A hardware 64-bit, 33MHz PCI interface to the SBC is implemented with a PLX 9656 PCI Master chip. The SBC, in the adapter, has its front panel at the face of the crate and is easily removable for upgrade or repair. The modular design provides a clear path for CPU performance upgrades by simple swapping of SBC cards.

[12] http://www.Advantech-nc.com.
[13] http://www.vmic.com
[14] http://www.dtims.com.
[15] http://www.teknor.com
[16] Tundra Semiconductor Corp., http://www.tundra.com.

Given comparable I/O capabilities, the amount of time required to run complex algorithms should be inversely proportional to the processor speed; more complicated algorithms can be used to process the available data if the processors are faster. However, an increase of processing power is more useful when supplied in the form of a single processor than in the form of a second identical processor working in parallel. This is because load balancing among multiple nodes is difficult in the Level 2 system due to constraints imposed by the front-end digitization. The front-end digitization holds buffers for 16 events awaiting Level 2 trigger decisions. A critical restriction in the system is that L2 results (accept or reject) must be reported to the front-end buffers in the order in which the triggers were taken at L1. While one processor works on an event with a long processing time, other events will arrive and fill the 16 front-end buffers. Other processors working on these events will go idle if they finish processing them quickly, since they cannot receive new events until the pending decision on the oldest event is taken. In other words a farm model is not appropriate for processing events at Level 2. Faster processing for each event in turn is thus more desirable than adding additional processors, once a baseline level of parallelism is established.

The quality of the Run IIb physics program will depend in large measure on effective rejection of background events in this more demanding environment. The Level 2 βeta upgrade will provide more resources needed to keep Level 2 in step with these demands and to further improve on background rejection from an upgraded Level 1. A subset of the most heavily-loaded processors should be replaced with higher-performance processors. Assuming that processors in the format used by the L2βetas increase performance by Moore's law, a purchase near the start of Run IIb could gain another factor of 4 in processing power over the first L2βeta processors.

6.3 Run IIb Algorithm Changes

We have begun to consider Run IIb algorithms that would profit from additional CPU power. Several performance enhancements are summarized in this section.

With longitudinal sectors on the order of 10cm vertex resolution of order ~few cm should be possible at L2 with the tracks output by the Level 2 Silicon Track Trigger (L2STT). At present all calorimeter tower information is reported to L2 relative to a vertex position of 0cm (centered with respect to the detector). The applications of vertex information at L2 are several.

To first order fast processors will supply resources to correct calorimetric information for vertex positions and significantly enhance resolutions at Level~2 for jets (including *tau*s), electro-magnetic objects, and Missing ET. Improved resolutions allow higher trigger thresholds for a given signal efficiency. For example, in the case of single jet triggers of moderate transverse energy, if 10cm vertex resolution is available, trigger rates may be improved by as much as 30-

60%, depending on pseudorapiditiy constraints[17]. This corresponds to a rejection factor of 1.4-2.5 independent of the L1 which, for the calorimeter, will absorb much of the Run IIa rejection power in Run IIb. The most straight forward model to implement these corrections is to add a track-based vertex finding algorithm to the Level 2 Central Tracking Trigger (L2CTT) preprocessor and to add a vertex correction flag to calorimeter objects processed in L2 Global. Preliminary timing measures for a track based vertexing algorithm in Level 3, show that such algorithms can find vertices in well under a millisecond (Figure 77). With advancement in CPU performance, further algorithm optimization, and parallelizing the algorithm between multiple processors, a similar algorithm will be tenable for STT tracks at Level 2.

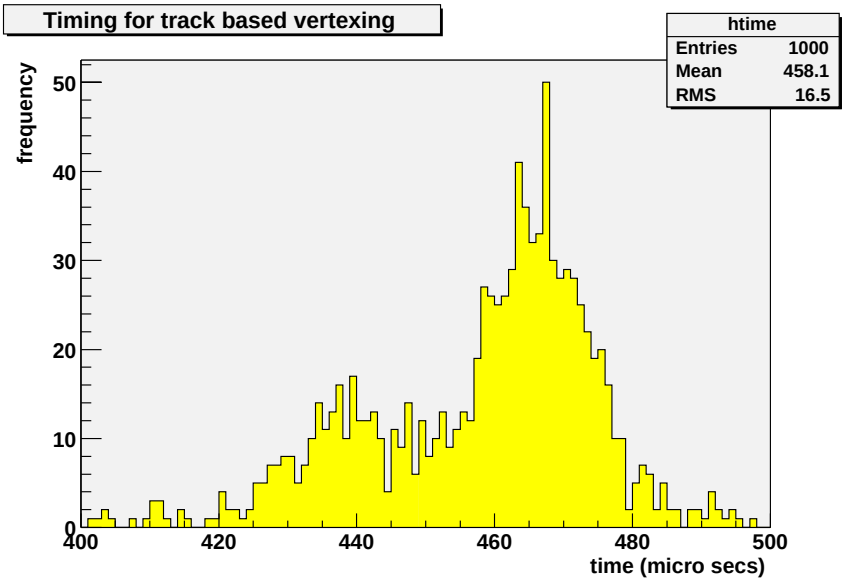

Figure 77. Timing distribution for vertex finding from L3 tracks (on a PIII 1.5 GHz cpu). Monte Carlo events used for this study were from a t-tbar sample with and average of 2.5 minimum bias interactions overlayed.

Another strategy to improve the resolution for calorimeter objects at L2, would be to apply tower by tower calibration corrections (or thresholds in the case of missing ET in the calorimeter processor). Such corrections would add linearly to the processing time by at least 35-70μs[18] on the 850MHz processors under study with the βeta prototypes. CPUs with three or more times this power of Run IIa can comfortably move these calculations within our nominal 50-100μs budget.

Multi-track displaced vertices could be searched for with the tracks output by the L2STT. This is beyond the original projected work of the Level 2 Central

[17] Level-2 Calorimeter Preprocessor Technical Design Report, http://hepalpha1.phy.uic.edu/l2cal/
[18] Time measured to apply correction factors to all calorimeter towers in software, the low limit is for correcting only total tower energies, the upper limit if for correcting EM and HADRONIC components separately.

Tracking Trigger preprocessor, and would be more CPU intensive. On another front, a sophisticated neural-net filter may search for tau events in the L2 Global processor. The effectiveness of such improvements depends on the actual mix of triggers chosen for Run IIb physics, so these should only be considered as examples. We have not yet studied which algorithms can be imported from Level 3 and applied to the lower-precision data available in Level 2.

A clear need for additional CPU power is in the global processor, which does the work of final Level 2 trigger selection by combining the results of preprocessors across detectors. More powerful CPUs will allow us to break the present software restriction of one to one mapping of Level 1 and Level 2 trigger bits (128 each at this point). This would allow more specific trigger processing to be applied to individual L1 trigger conditions at Level 2, as we currently do in Level 3. In addition to channels with inherent physics interest, many signals will play increasingly important roles in the calibration of the detector and efficiency measures for the main physics menu's selection criteria. Added trigger branching will greatly facilitate the collection of these data. It is at times impossible to simulate a dataset with the necessary accuracy to calculate efficiencies and acceptances for complex trigger conditions, especially when hardware calibration effects have exceedingly strong bearing. The cpu cost in trigger branching would grow at least linearly with the number of triggers reported.

The anticipated distribution of the new processors is as follows:

- Calorimeter (2 SBCs) - Apply data corrections to improve ET resolution jets, electrons, Missing ET.
- Global (3 SBCs) – Apply vertex corrections to calorimeter objects, improve b-tagging by searching for multi-track displaced vertices. Enhanced trigger branching.
- Tracker (2 SBCs) – Handle increased number of silicon layers, calculate quantities needed for z-vertex and multi-track displaced vertices.Muon (1 SBC) - Maintain rejection at high occupancy.Preshower (2 SBCs) – Maintain rejection at high occupancy.
- Spare/test (2 SBCs) – Spare+“shadow” nodes for test/development purposes.

In addition to the primary upgrade path of adding higher power CPU cards, a further upgrade avenue may include equipping the cards with dual processors that share the card’s memory and I/O. This upgrade is attractive because its incremental cost is low, but it will require a substantial software effort to turn it into increased throughput, even if it is possible to build code that takes advantage of the dual processors without writing thread-safe code. However, a dual-processor upgrade might be attractive for reasons other than performance. One processor could keep the Linux operating system active for debugging of problems in algorithms run in the second processor. Or one could run a production algorithm in one processor and a developmental version in the

second processor. This second processor might even be operated in a “shadow” mode (as in Level 3), processing events parasitically, but skipping events if the developmental algorithm gets behind, or is being debugged. These possibilities will be studied prior to Run IIb, though dual CPU cards are not intended as a substitute for higher power upgrade processors.

6.4 Summary

For Run IIb, we are proposing a partial upgrade of the Level 2β system that replaces the processors on 12 boards. This is in anticipation of the potential increase in computing power that could at that time be used to implement more sophisticated tracking, STT, and calorimeter/track matching algorithms at Level 2 in response to the increased luminosity.

7 Level 2 Silicon Track Trigger

7.1 Motivation

The DØ Level 2 Silicon Track Trigger (L2STT) receives the raw data from the SMT on every level 1 accept. It processes the data from the axial strips in the barrel detectors to find hits in the SMT that match tracks found by the level 1 track trigger in the CFT. It then fits a trajectory to the CFT and SMT hits. This improves the resolution in momentum and impact parameter, and the rejection of fake tracks, compared to the central track trigger alone.

The L2STT matched to the Run IIa SMT detector is being constructed with NSF and DOE funds for delivery in 2002. An upgrade for Run IIb, however, will be necessary in order to match the new geometry of the Run IIb Silicon Tracker. Upgrading the L2STT to optimize its rejection power by using all of the information from the new Run IIb SMT is an important part of maintaining the rejection of the Level 2 trigger in Run IIb.

Tracks with large impact parameter are indicative of long-lived particles (such as b-quarks) which travel for several millimeters before they decay. The L2STT thus provides a tool to trigger on events with b-quarks in the level 2 trigger. Such events are of particular importance for the physics goals of Run 2. The Higgs boson decays predominantly to $b\bar{b}$ pairs if its mass is less than about 135 GeV/c^2. The most promising process for detection of a Higgs boson in this mass range at the Tevatron is associated production of Higgs bosons with W or Z bosons. If the Z boson decays to neutrino pairs, the b-quarks from the Higgs decay are the only detectable particles. In order to trigger on such events (which constitute a significant fraction of associated Higgs production) the L2STT is essential to detect at the trigger level jets that originate from b-quarks. The L2STT will also allow the collection of a large enough sample of inclusive $b\bar{b}$ events to see the decay $Z \rightarrow b\bar{b}$. Such a sample is important to understand the mass resolution and detection efficiency for $b\bar{b}$ resonances, and to calibrate the calorimeter response to b-quark jets. The latter will also help to drastically reduce the uncertainty in the top quark mass measurement, which is dominated by the jet energy scale uncertainty. Detailed descriptions of the physics benefits of STT are written up as DØ Notes[19,20].

7.2 Brief description of Run IIa STT architecture

The STT is a level-2 trigger preprocessor, which receives inputs from the level 1 central track trigger (L1CTT) and the silicon microstrip tracker (SMT). The STT filters the signals from the SMT to select hits that are consistent with tracks found by L1CTT. The L1CTT uses only the axial fibers of the CFT to find track patterns. No z-information is available for level-1 tracks and SMT hits are filtered

[19] "A silicon track trigger for the DØ experiment in Run II – Technical Design Report", Evans, Heintz, Heuring, Hobbs, Johnson, Mani, Narain, Stichelbaut, and Wahl, DØ Note 3510.
[20] "A silicon track trigger for the DØ experiment in Run II – Proposal to Fermilab", DØ Collaboration, DØ Note 3516.

based only on their r-ϕ coordinates. Then the L2STT fits a trajectory to each level-1 track and the associated selected hits. In the fit, only axial information is used. Matching axial and stereo hits from the SMT is too complex a task to complete in the available time budget. In the selection of the SMT hits, however, the constraint is imposed that they originate from at most two adjacent barrel sections. The distribution of the hit pattern over the two barrel-sections must be consistent with a track. The fit improves the precision of the measurements of transverse momentum and impact parameter, compared to the level-1 track trigger. It also helps reject fake level-1 tracks for which there are no matching SMT hits.

The STT processes these data for 12 azimuthal sectors independently. Each sector consists of 36 detector elements in four radial layers and six barrel segments. The geometry of the SMT in Run IIa provides enough overlap between adjacent detector elements that each detector element can be uniquely associated with one of these sectors without significant loss of acceptance due to tracks that cross sectors.

There are three distinct functional modules in the STT. The fiber road card (FRC) receives the data from L1CTT and fans them out to all other cards that process hits from the same sector. The silicon trigger card (STC) receives the raw data from the SMT front ends and filters the hits to associate them with level-1 tracks. The track fit card (TFC) finally fits trajectories to level-1 tracks and SMT hits. Each of these modules is implemented as a 9Ux400 mm VME card, based on a common motherboard. The main functionality is concentrated in large daughter boards, which are distinct for the three modules. Communication between modules is achieved through serial links. The serial links use low voltage differential signaling (LVDS) at 132 MB/s. We designed PC-MIP standard mezzanine boards that accommodate either 3 LVDS transmitters or 3 LVDS receivers. The motherboard has six slots to accommodate these boards. Data are transferred between the VME bus, the daughter cards and the link mezzanine boards over three interconnected 32-bit/33 MHz PCI busses. Figure 78 shows a block diagram and photograph of the motherboard.

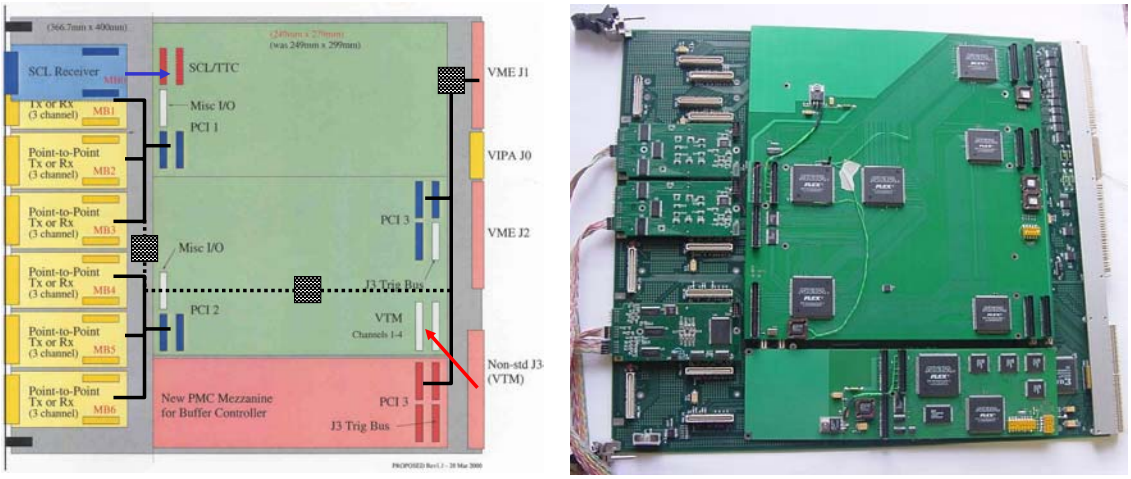

Figure 78. Block diagram and photograph of motherboard. The drawing shows the three PCI busses (solid lines) and the bridges that connect them (squares and dashed lines). The photograph also shows the FRC daughter board, the buffer controller mezzanine board, two LTBs, and one LRB.

The FRC module receives the data from the L1CTT via one optical fiber and a VTM in the rear card cage of the crate. The FRC also receives information from the trigger control computer via the serial command link (SCL). This information contains the level-1 and level-2 trigger information and identifies monitor events for which all the monitor counters have to be read out. The FRC combines the trigger information with the road data and sends it to all other modules in the crate via serial links. The motherboard can accommodate up to six PC-MIP mezzanine boards. One is used to receive the SCL; the remaining five can be used for LVDS transmitter boards to fan out the L1CTT data, which provides up to 15 links. The FRC also performs arbitration and control functions that direct the flow of data for accepted events to the data acquisition system. The buffer controller mezzanine board (BC) holds a multiport memory in which events are stored until a level-2 trigger decision has been taken. There is one BC on each motherboard. The FRC manages the buffers on all BCs in the crate.

The STC module receives the L1CTT data from the FRC over an LVDS serial link. Each STC module has eight channels, which each process the data from one silicon detector element. The signals from the SMT front ends are transmitted over a 106 MB/s serial link using the HP G-link chips and optical fibers from the electronics platform below the detector to the 2nd floor of the moveable counting house, where the STT is located. Passive optical splitters create two data paths, one to the SVX data acquisition system and another into the STT. The optical signals are received and parallelized in VME transition modules (VTM) sitting in the rear card cage of the crates that accommodate the STT modules. The VTMs are an existing Fermilab design, used by both DØ and CDF. The SMT signals are passed through the J3 backplane to the STC module sitting in the main card cage in the same slot as the VTM. Each VTM has four optical receivers and each fiber carries the signals from two detector elements.

In the STC module, each level-1 track is translated to a range of strips (a “road”) in each of the eight detector elements that may contain hits from the particle that gave rise to the level-1 track using a look-up table. The SMT data are clustered to combine adjacent strips hit by the same particle. These hits are then compared to the roads defined by the level-1 tracks. The hits that are in one or more roads are queued for transfer to the TFC module over an LVDS serial link. The main logic of the STC module is implemented in a single large field programmable gate array (FPGA).

Each TFC receives all hits from one azimuthal sector that were associated with at least one road. Because of the way SMT detector elements are mapped onto the optical fibers, three STC modules receive hits from both sectors in the crate. The outputs of these three STC modules go to both TFC modules in the crate. The remaining six STC modules receive hits from only one sector and their outputs go to only one TFC module. Thus each TFC module has six incoming LVDS serial links. The hits that come in over these links are sorted according to the level-1 track they are associated with. Then all data associated with one level-1 track is sent to one of eight DSPs that perform a linearized chi-squared fit. The results of the fits and the L1CTT data are sent via a Cypress hotlink to the level-2 central track trigger (L2CTT). The L2CTT acts as a concentrator for the 12 hotlink inputs from the six STT crates.

The number of crates required for the entire system is driven by the number of STC modules required to instrument all barrel detectors. Each SMT module can process the data from eight detector elements. Each azimuthal sector consists of 36 detector elements. Thus, each azimuthal sector requires 4.5 STC modules. We can accommodate two such sectors in one VME crate, so that one crate contains one FRC module, nine STC modules, and 2 TFC modules (one per azimuthal sector). In addition, each STT crate also houses a power PC and a Single Board Computer (SBC) card. The former controls the VME bus, and is used to download data tables and firmware into the STT modules and to monitor the performance of the STT. The SBC transfers the data of accepted events to the data acquisition system. Figure 79 shows the layout of one STT crate.

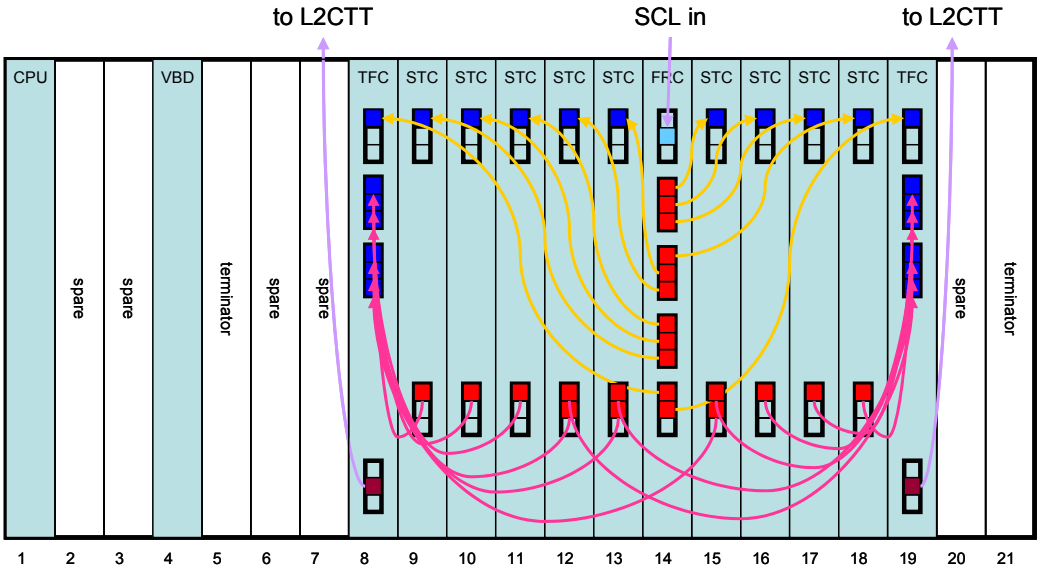

Figure 79. Layout of L2STT crate for Run IIa. The groups of three squares on the front panels indicate PC-MIP boards and the colored squares indicate used channels. The light blue square at the top of the FRC indicates the SCL receiver, and the brown squares at the bottom of the TFCs indicate the hotlink transmitters. Arrows indicate cable connections and are directed from LTBs (red) to LRBs (blue).

7.3 Changes in tracker geometry and implications for STT

The design of the silicon microstrip tracker for Run IIb[21] foresees six concentric layers of detector elements, compared to four for the Run IIa design. The inner two layers consist of twelve 78 mm long sensors along the beam direction. Layers 2 and 3 consist of ten 100-mm long sensors and the outermost layers consist of twelve 100-mm long sensors. Figure 80 shows two views of the design.

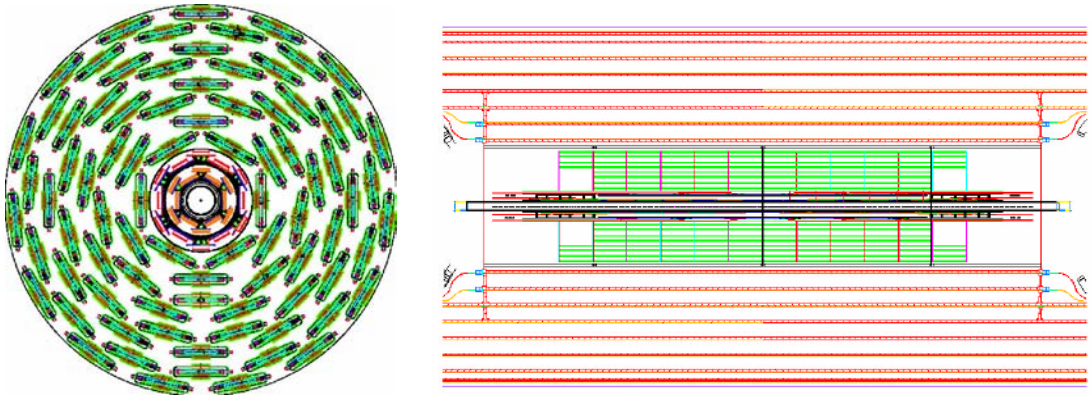

Figure 80. Axial and plan views of the Run IIb silicon microstrip tracker design.

[21] "DØ Run IIb Silicon Detector Upgrade - Technical Design Report", DØ Collaboration, 2001.

Some sensors are ganged and in layers 1-5 every cable reads out the signals from two sensors to reduce the number of readout units (i.e. cables). Table 39 lists the number of readout units for axial strips in every layer, which determines the number of STC modules required to process their hits in the STT. The detector elements in each layer alternate between the two radii listed in the table such that adjacent detectors overlap slightly. Readout units for stereo strips are not used in the STT and are therefore not listed here. The number of readout units with axial strips increases from 432 in the Run IIa design to 552 in the Run IIb design.

Table 39 Parameters of Run IIb silicon microstrip tracker design.

Layer	Radius (axial strips)	Strip pitch	Strips	Readout units in ϕ	Readout units in z
0	18.6/24.8 mm	50 μm	256	12	12
1	34.8/39.0 mm	58 μm	384	12	6
2	53.2/68.9 mm	60 μm	640	12	4
3	89.3/103 mm	60 μm	640	18	4
4	117/131 mm	60 μm	640	24	4
5	150/164 mm	60 μm	640	30	4

The data must be channeled into TFCs such that all hits from a track are contained in one TFC. In layers 0, 1, and 2 the overlaps between adjacent detector elements are large enough so that each sensor can be uniquely associated with one TFC. This divides the detector into 12 azimuthal sectors as indicated by the shaded regions in Figure 81. To maintain full acceptance for tracks with p_T>1.5 GeV/c and impact parameter b<2 mm, the data from some sensors in layers 3, 4, and 5 must be channeled into two TFCs, which are in some cases located in different crates. This is not the case in the current configuration, but should not present any problems. We are limited to 8 STC inputs into each TFC, which is sufficient for the Run IIb detector geometry.

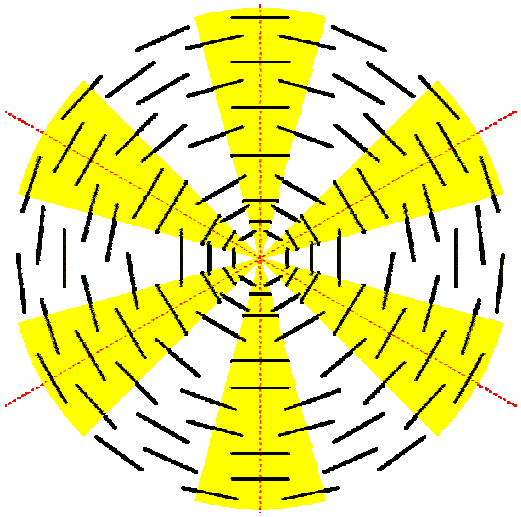

Figure 81. Azimuthal sector structure in the Run IIb silicon microstrip tracker. The yellow wedges indicate the coverage of the 12 azimuthal sectors. All detector elements in layers 3-5, which cover parts of two sectors are assigned to two TFC modules.

The hardware of the Run IIa STT is only sufficient to instrument four of the six layers of the Run IIb SMT. To include five or all six of the layers in the trigger, additional modules need to be acquired. For the most part this amounts to building additional copies of the Run IIa STT modules. None or very little new hardware design is required. The following section explains the upgrade options that were considered in detail.

7.4 Simulation of the Run IIb STT

We have simulated the expected performance of various upgrade options for the STT under Run IIb conditions. We have used simulated event samples generated with the PYTHIA event generator and a GEANT simulation of the Run IIb silicon microstrip tracker (SMT) geometry and the central fiber tracker (CFT). The signal we are interested in triggering on is the production of b-quarks in the decay of a heavy particle. This process is represented by a sample of events from the process $p\bar{p} \to WH$, followed by $W \to \mu\nu$ and $H \to b\bar{b}$. The backgrounds we want to discriminate against are light quark and gluon jets. This is represented by a sample of events from the process $p\bar{p} \to Z \to q\bar{q}$, where $q\bar{q}$ is a pair of light quarks (u or d).

At the Run IIb luminosity of $5\text{x}10^{32}$ $\text{cm}^{-2}\text{s}^{-1}$ and a crossing interval of 132 ns, we expect to see on average about five soft proton-antiproton interactions in every triggered hard scattering event. We simulate the additional interactions with minimum bias events generated using PYTHIA. Comparison with data from Run 1 show that PYTHIA underestimates the particle multiplicity in these additional interactions so that we need on average 7.5 additional PYTHIA minimum bias events to simulate the conditions expected for Run IIb. We therefore superimpose additional interactions on the above events with a Poisson distribution with a mean of 7.5 (N_{mb}=7.5). To illustrate the effect of these

additional interactions, we also show the performance of the level 2 silicon track trigger for simulated event samples without additional interactions (N_{mb}=0).

In order to simulate the level 1 track trigger inputs, we have used the tracks found in the CFT data by a track finding program. To simulate the limited track information used in the STT, we use only the hits closest to these tracks in the innermost (A) and outermost (H) layers of the CFT. The SMT inputs are clusters of hits in detectors with axial silicon microstrips.

The CFT tracks define roads that go through the interaction point at the center of the detector. In each of the six SMT layers the cluster is used that is closest to the center of the road. A linearized trajectory given by $\phi(r) = b/r + kr + \phi_0$ is fit to the hits in the six silicon layers and the A and H layers of the CFT. The parameter r is the distance in the xy plane from the interaction point, b is the impact parameter, k the radius of curvature, and ϕ_0 the azimuth of the track. If $\chi^2/dof > 5$, the hit with the largest contribution to χ^2 is dropped and the trajectory refit. We require that at most one silicon layer is missing a hit and that $\chi^2/dof < 5$ for all good STT tracks. We find that the impact parameter resolution is about 11 μm for good high-p_T tracks. The resolution gets worse for tracks with lower p_T (39 μm at 1 GeV) and with missing or dropped hits in the fit. We parameterize the impact parameter resolution σ as a function of these parameters and define the impact parameter significance $s_b = b/\sigma$. Figure 82 (a) shows the distribution of s_b for tracks from the WH and Z samples with N_{mb}=0. The figure shows that tracks from displaced vertices are responsible for the tails of the impact parameter significance distribution and that the tracks from the Z sample, which should all be prompt, are almost Gaussian in their distribution, as expected.

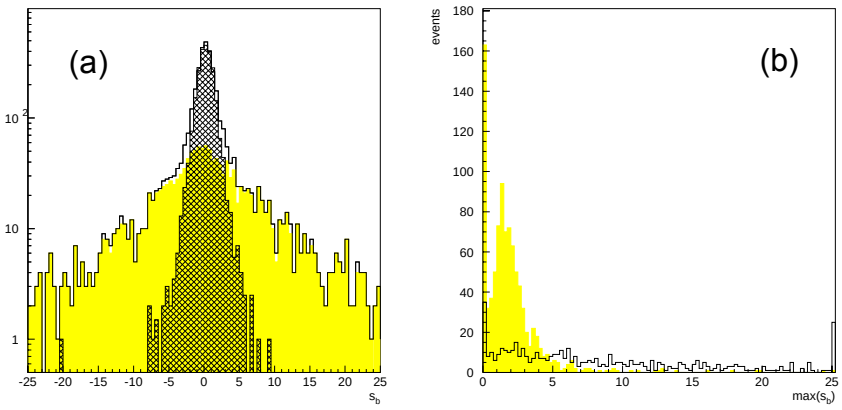

Figure 82: (a) Distribution of impact parameter significance of good tracks from the WH and Z samples with N_{mb}=0. The open histogram shows all good tracks from the WH sample, the colored histogram shows the subset that is matched to Monte Carlo particles from displaced vertices and the hatched histogram shows all good tracks from the Z sample. (b) Distribution of the largest

value of impact parameter significance per event for good tracks with p_T>1.5 GeV. The colored histogram shows the Z sample and the open histogram the WH sample, both with N_{mb}=0.

We trigger on an event if there is a good STT track with s_b greater than some threshold. Figure 82 (b) shows the distribution of the largest impact parameter significance for good STT tracks per event. The trigger efficiency is given by the number of WH events that have a good STT track with s_b greater than a threshold. The rejection is the inverse of the efficiency for the Z event sample.

Figure 83 shows the rejection versus efficiency curves from event samples with N_{mb}=0 and 7.5 using all six silicon layers. We see that the rejection at fixed efficiency drops by about a factor 2 due to the additional events. We then remove silicon layers from the STT processing. We considered removing layer 4, layer 0, and layers 1 and 3 together. Rejection at fixed efficiency drops every time a layer is removed. Removing layer 4 reduces the rejection by about 20%. Removing layer 0 reduces the rejection by about a factor 2, as does removing layers 1 and 3 together. We tabulate some benchmark values in Table 40.

Table 40: Benchmark values for rejection achieved by STT for different conditions.

SMT layers used	*N_{mb}*	*rejection for 65% efficiency*
012345	0	22
012345	7.5	11
01235	7.5	9
12345	7.5	6
0245	7.5	6

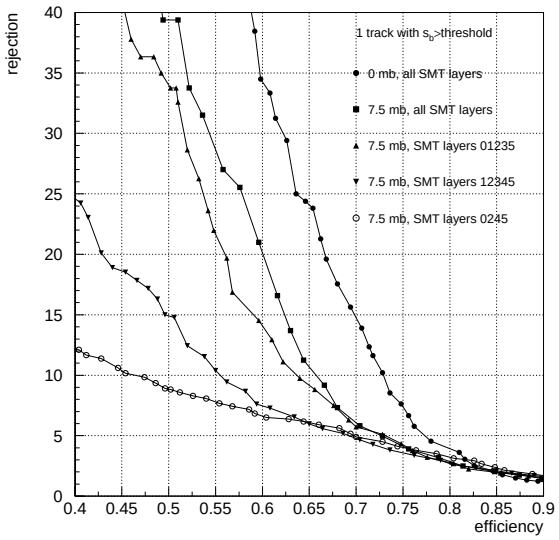

Figure 83: Curves of rejection versus efficiency for triggering on one good STT track with impact parameter significance above a threshold and p_T>1.5 GeV. The curve marked with ● is for N_{mb}=0 and using all six silicon layers. All other curves are for N_{mb}=7.5 and for various combination of silicon layers included in the trigger.

Figure 84 shows curves of rejection versus efficiency when the p_T threshold for the CFT tracks that define the roads is varied. In Run IIa, the level 1 track trigger can detect tracks with p_T>1.5 GeV. As the figure shows, it is important to maintain this capability in Run IIb, since the rejection at fixed efficiency drops when this threshold is raised above 1.5 GeV.

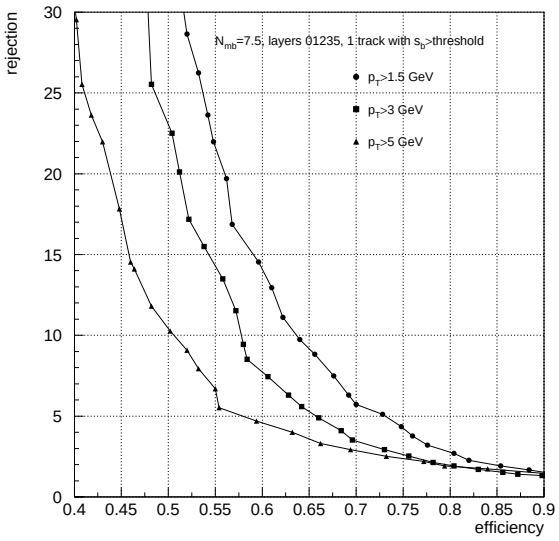

Figure 84: Curves of rejection versus efficiency for triggering on one good STT track with impact parameter significance above a threshold. For the three curves the minimum p_T of the CFT tracks that define the roads is varied as shown.

Aside from triggering on tracks from displaced vertices, the STT also helps reject fake level 1 track trigger candidates that are due to the overlap of other, softer, tracks. We find that at N_{mb}=7.5 and a track p_T threshold of 5 GeV, the STT only accepts 1 in 3.2 fake CFT tracks if all six silicon layers are used. This drops to 1 in 2.8 for 5 layers used and about 1 in 1.8 for 4 layers used in STT. This rejection of fake L1 track triggers is crucial since the rate of level 1 track triggers is expected to increase significantly at high luminosities as shown in section 3.

In Run IIb, the processing times and latencies for the STT preprocessor will become larger. The additional hits from the higher luminosity and the additional silicon detectors will increase transfer times. The timing of the STT is dominated by the fitting process in the TFCs. There are two main components that determine this timing. About 0.6 μs per hit are required to select the hit to include in the fit for each layer. This time scales linearly with the number of hits per road. The actual fit takes about 7-14 μs per track, depending on whether the fit is repeated after dropping a hit. Both components will increase in Run IIb.

Figure 85 shows the number of clusters in a 2 mm wide road for WH events with N_{mb}=7.5. If layers 0, 1, 2, 3, and 5 are used in the trigger, we expect on average 26 hits in the road of a given track. This is larger than in Run IIa, because of the closer distance to the interaction point of the innermost silicon layer and because of the larger number of layers.

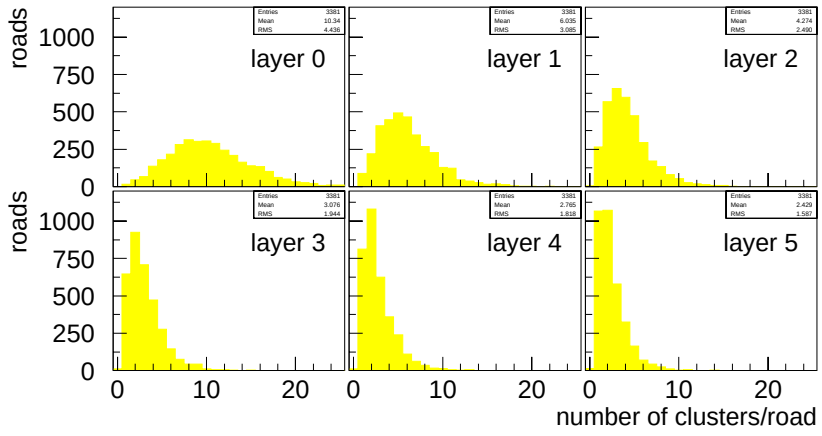

Figure 85: Distributions of hit multiplicity per road per layer for the WH sample with N_{mb}=7.5.

The time required for the fit will increase because of the additional layer included in the fit relative to Run IIa. In addition, for the large multiplicity in the first two layers, our hit selection algorithm may be inadequate. The algorithm currently selects the hit that is closest to the center of the road thus biasing the fit towards small impact parameters. We have in the past investigated different algorithms in which the road is not forced to the interaction point. These require more processing time and were not required for the lower luminosities of Run IIa.

Queuing simulations show that the STT operates within its time budget of about 100 μs on average. However latencies up to 250 μs are observed and these times will increase in Run IIb for the reasons mentioned in the previous paragraphs. In order to avoid exceeding our time budget, we will require additional processor cards (TFC).

From this work, we conclude that the STT must be upgraded to include at least five of the six layers of the Run IIb SMT. Without any upgrade, the STT performance will be severely degraded. Ideally, we would like to instrument all six silicon layers to create a trigger with the most rejection power. Considering the fiscal constraints, however, we propose to upgrade the STT to instrument five silicon layers (0, 1, 2, 3, and 5). Since we likely to exceed our time budget with the increased processing time required for the additional silicon layer, we also propose to double the number of TFCs so that two TFCs are assigned to each azimuthal sector.

7.5 Implementation description for STT upgrade

The layout for an STT crate during Run IIb is shown in Figure 86.

Instrumenting SMT layers 0, 1, 2, 3, and 5 requires one additional STC and the associated VTM per crate. Increasing the CPU power for fitting requires two additional TFCs per crate. All of these will require motherboards. We have to build additional link transmitters and receivers to ship the data to and from the new boards. The designs for all of these already exist in the Run IIa STT. Thus we have to only go through additional production runs. Additional optical splitters and fibers must also be purchased for the larger number of silicon input channels.

In order to maintain the same number of output cables into L2CTT we have to merge the outputs of the two TFCs that are assigned to the same channel. We will achieve this by daisy-chaining the two TFCs. That means one TFC uses the existing hotlink transmitter to send its output into the second TFC. It will be received by a hotlink repeater, which will have all the functionality of the existing hotlink transmitter and in addition a hotlink receiver and logic to merge the two data streams before they are sent to L2CTT.

There are three spare slots in the J3 backplane that can accommodate the new boards. Thus no modification has to be made to the crates. Our power supplies are also dimensioned large enough for the additional boards.

The fitting algorithm in the TFCs has to be modified to reflect the different number of layers and new coordinate conversion tables will have to be computed. The firmware in the STCs will have to be modified to handle the modified inputs from the L1CTT and new road look-up tables will have to be computed.

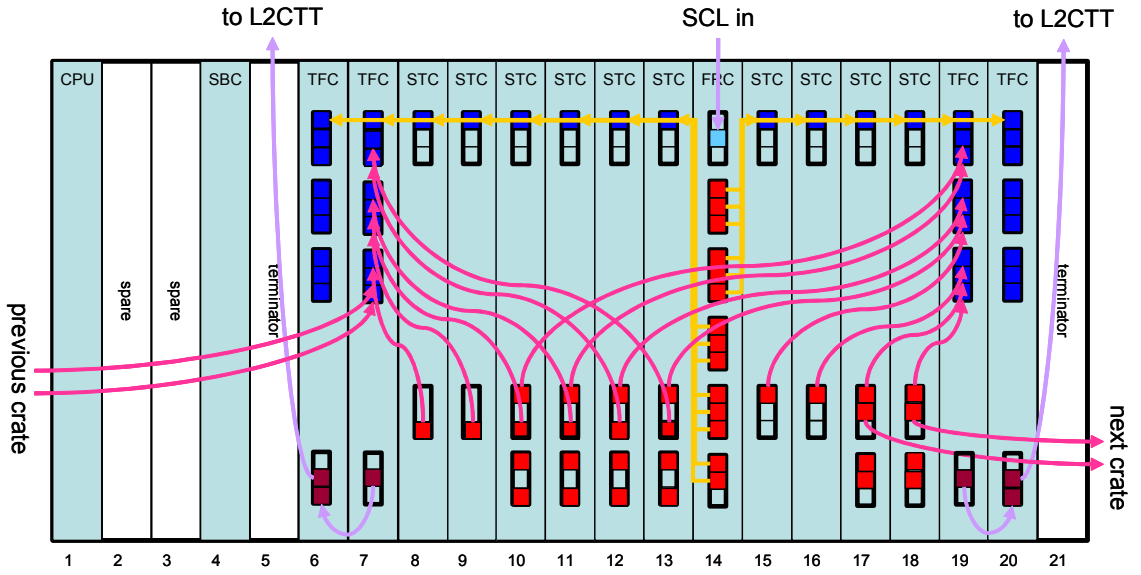

Figure 86. Layout of STT crate for Run IIb. The groups of three squares on the front panels indicate PC-MIP boards and the colored squares indicate used channels. The light blue square at the top of the FRC indicates the SCL receiver, and the brown squares at the bottom of the TFCs indicate the hotlink transmitters and repeaters. Arrows indicate cable connections and are directed from LTBs (red) to LRBs (blue). For clarity, the cables from STCs to TFCs are shown only for two of the four TFCs.

8 Trigger Upgrade Summary and Conclusions

The DØ experiment has an extraordinary opportunity for discovering new physics, either through direct detection or precision measurement of SM parameters. An essential ingredient in exploiting this opportunity is a powerful and flexible trigger that will enable us to efficiently record the data samples required to perform this physics. Some of these samples, such as $p\bar{p} \to ZH \to b\bar{b}\,\nu\bar{\nu}$, are quite challenging to trigger on. Furthermore, the increased luminosity and higher occupancy expected in Run IIb require substantial increases in trigger rejection, since hardware constraints prevent us from increasing our L1 and L2 trigger rates. Upgrades to the present trigger are essential if we are to have confidence in our ability to meet the Run IIb physics goals.

To determine how best to meet our Run IIb trigger goals, a Run IIb Trigger Task Force was formed to study the performance of the current trigger and investigate options for upgrading the trigger. Based on the task force recommendations, we have adopted the following plan for the trigger upgrade:

1. Replacement of the Level 1 Central Track Trigger (CTT) DFEA daughter boards. The CTT is very sensitive to occupancy in the fiber tracker, leading to a large increase in the rate for fake high-p_T tracks in the Run IIb environment. The new daughter board will utilize more powerful FPGAs to implement individual fiber "singlets" in the trigger, rather than the "doublets" currently used. Preliminary studies show significant reductions in the rate of fake tracks can be achieved with this upgrade.
2. Replacement of the Level 1 calorimeter trigger. The calorimeter trigger is an essential ingredient for the majority of DØ triggers, and limitations in the current calorimeter trigger, which is essentially unchanged from the Run 1, pose a serious threat to the Run IIb physics program. The two most serious issues are the long pulse width of the trigger pickoff signals and the absence of clustering in the jet trigger. The trigger pickoff signals are significantly longer than 132 ns, jeopardizing our ability to trigger on the correct beam crossing. The lack of clustering in the jet trigger makes the trigger very sensitive to jet fluctuations, leading to a large loss in rejection for a given trigger efficiency and a very slow turn-on. Other limitations include exclusion of ICD energies, inability to impose isolation or HAD/EM requirements on EM triggers, and very limited capabilities for matching tracking and calorimeter information. The new L1 calorimeter trigger would provide:
    - A digital filter that utilizes several samplings of the trigger pickoff signals to properly assign energy deposits to the correct beam crossing.
    - Jet triggers that utilize a sliding window algorithm to cluster calorimeter energies and significantly sharpen jet energy thresholds.
    - Inclusion of ICD energy in the global energy sums to improve missing E_T resolution.

- Electron/photon triggers with the ability to impose isolation and/or HAD/EM requirements to improve jet rejection.
- Topological triggers that aid in specific event topologies, such as acoplanar jets.

3. A new calorimeter-track match system. Significant improvements in rates have been demonstrated for both EM and track-based τ triggers from correlating calorimeter and tracking information. The cal-track match system utilizes boards that have already been developed for the muon-track matching system.
4. No major changes are foreseen for the Level 1 Muon trigger. Since the muon trigger matches muon and tracking information, it will benefit indirectly from the track trigger upgrade.
5. Some of the L2βeta processors will be replaced to provide additional processing power.
6. The L2 Silicon Track Trigger (STT) requires additional cards to accommodate the increased number of inputs coming from the Run IIb silicon tracker.
7. Maintaining Level 3 trigger rejection as the luminosity increases will require increasing the processing power of the L3 processor farm as part of the upgrade to the online system (see Part IV: DAQ/Online Computing).

Simulation studies indicate that the above upgrades will provide the required rejection for the Run IIb physics program. In particular, the expected trigger rate for the primary Higgs channels, WH and ZH, is compatible with our trigger rate limitations. The technical designs for these systems are making rapid progress. The designs are largely based on existing technologies, such as FPGAs and commercial processors, minimizing the technical risk. We foresee no major technical hurdles in implementing the proposed trigger upgrade.

(This page intentionally left blank)

DØ Run IIb Upgrade Technical Design Report

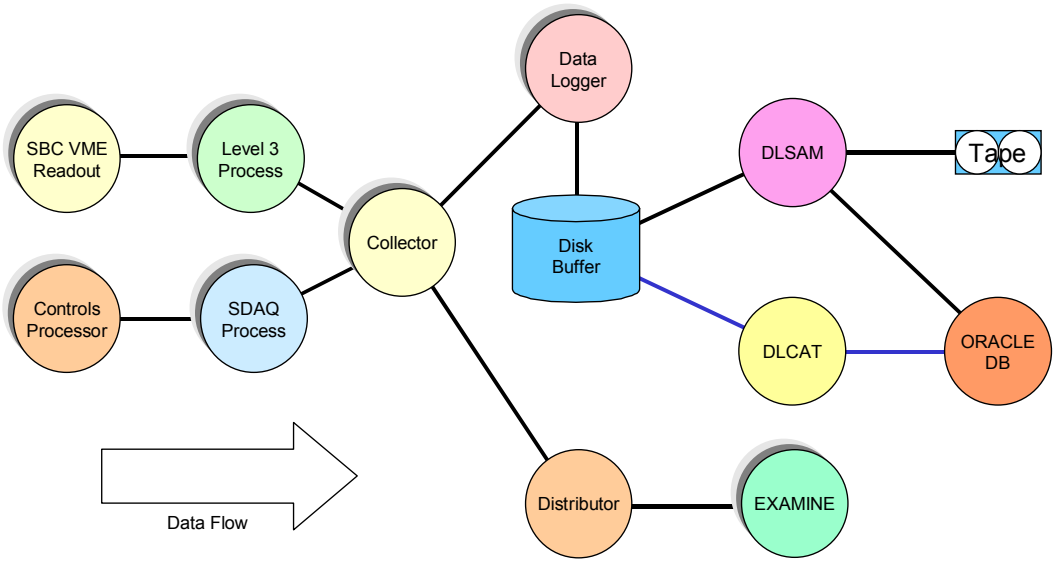

DAQ/ONLINE COMPUTING

(This page intentionally left blank)

DAQ/ONLINE COMPUTING CONTENTS

1 DAQ/Online Computing Overview

1.1 Scope

For the purposes of this document, the DØ DAQ and Online Computing systems will be defined to consist of the following components:

- DAQ and Online network,
- Single Board Computers (SBCs) in VME readout crates,
- Level 3 Linux software filter farm,
- Host Online system,
- Control room computing systems,
- Data monitoring computing systems,
- Database servers,
- File servers,
- Slow control system, including VME processors,
- plus the associated software for each of these elements.

1.2 Software Architecture

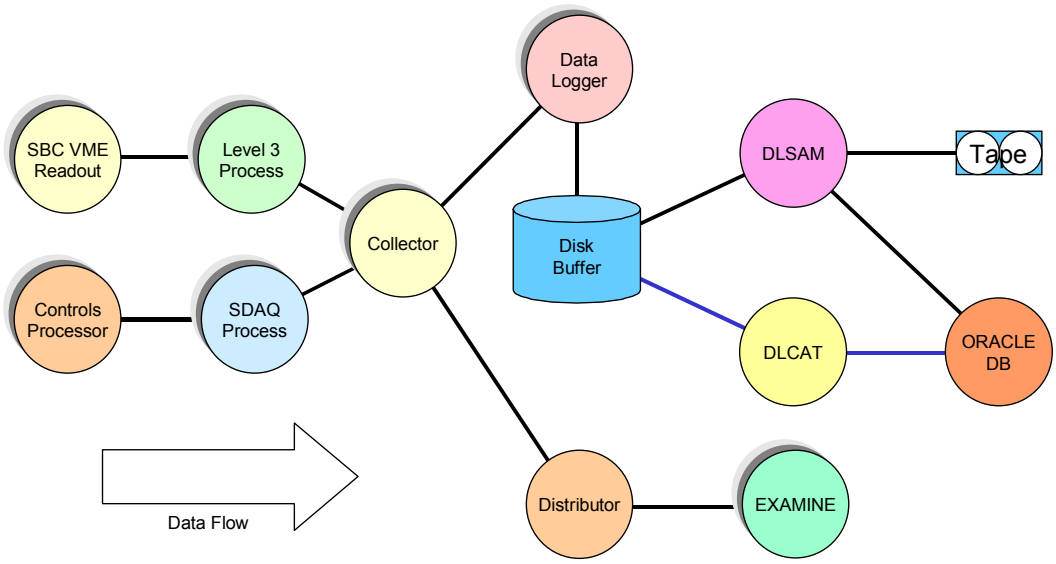

Figure 1. DAQ system software components

The software architecture of the Run IIb Online system is unchanged from that of Run IIa. Some components will need updating, but there are no structural differences planned. The major software components in the event data path are illustrated in Figure 1. The slow control system components are not illustrated in the figure, nor are the non-event monitoring systems.

1.3 Hardware Architecture

The hardware architecture of the Run IIb Online system is also unchanged from that of Run IIa. The current architecture is illustrated in Figure 2 and Figure

3. At the center of the system are two high capacity network switches (Cisco 6509). The Primary Data Acquisition (PDAQ) event data path includes the Single Board Computers in the VME readout crates, the Level 3 Linux filter nodes, the Host Online systems on which reside the Collector, Distributor, and Data Logger processes, and the final data repository in the Feynman Computing Center (FCC). The EXAMINE processes on the Monitor system nodes provide real-time event data analysis and monitoring functions. Some of the Slow Control system nodes also participate in the Secondary Data Acquisition (SDAQ) path that also feeds into the Host processes. An ORACLE database serves for configuration control and recording of run parameters. Also included in these figures are the Control Room, File Server, and Slow Control system nodes.

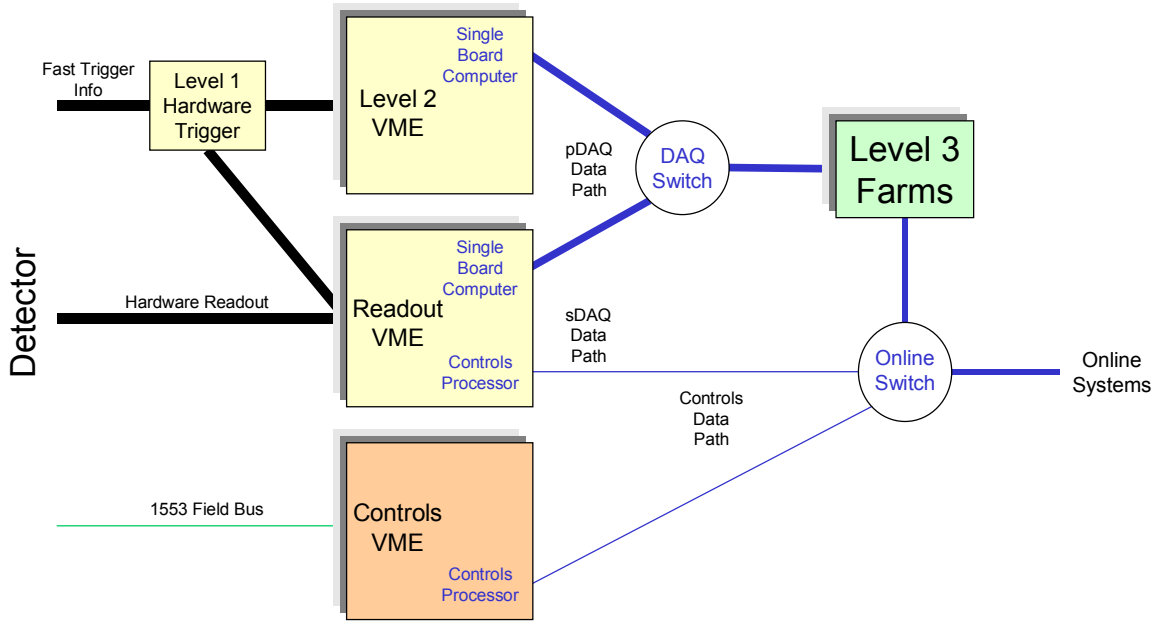

Figure 2. DAQ system hardware components.

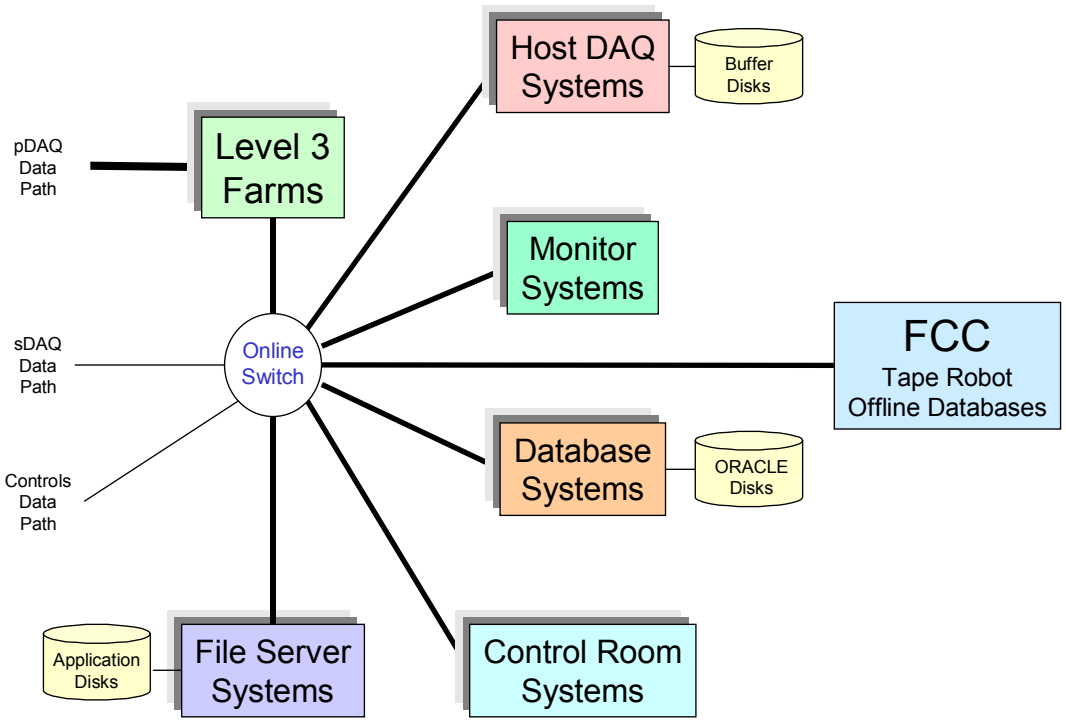

Figure 3. Online system hardware components.

For Run IIb many components of these computer systems will need to be updated or replaced.

1.4 Motivations

The primary considerations governing the development of the DØ Online system for Run IIb are supplying the enhanced capabilities required for this running period, providing hardware and software maintenance for the (by then) five-year old hardware, and supplying the required software support. We expect the requirements for the Online data throughput to at least double, largely driven by the ability of the Offline analysis systems to absorb and analyze the data. Many of the existing Online computing systems will reach the end of their viable lifetime in capability, maintainability, and software support by the Run IIb era. The gradual replacement of many of these component systems will be essential.

1.4.1 Enhanced Capabilities

The rate at which DØ records data to tape has been limited by the cost of storage media and the capability of the Offline systems to analyze the data. Assuming five years of improvements in computing capability, it is reasonable to expect that the Offline capacity for absorbing and analyzing data will more than double. The Online system must be capable of providing equivalent increased data throughput.

To analyze the higher occupancy events in the high-luminosity Run IIb era, more sophisticated software filters must run on the Level 3 trigger farm. These more complicated codes will increase execution time. The resulting increased

computing demand in Level 3 will need to be met by either an increase in the number of processors, replacement of these units by more capable processors, or both.

It is also expected that data quality monitoring software will be vastly improved by the Run IIb era. These capabilities again are likely to come at the cost of increased execution time and/or higher statistical sampling requirements. In either case, more numerous and more powerful monitoring systems will be required.

1.4.2 Hardware and Software Maintenance

By the time of Run IIb, the computing systems purchased for Run IIa will be more than five years old. In the world of computing hardware, this is ancient. Hardware maintenance of such old equipment is likely to be either impossible or unreasonably expensive. Experience shows that replacement by new (and under warranty) equipment is more cost effective. Since replacement of obsolete equipment not only addresses the maintenance question, but also issues of increased capability, it is the most effective course of action.

The DØ Online system is composed of several subsystems that have differing hardware components and differing maintenance needs. Subsystem specific issues will be addressed in the following sections.

1.4.3 Software Support

Several different operating systems are present in the Online system, with numerous custom applications. We have tried, wherever possible, to develop software in as general a fashion as possible so that it can be migrated from machine to machine and from platform to platform. However, support of certain applications is closely tied to the operating system on which the applications run. In particular, ORACLE database operations require expertise that is often specialized to the host operating system. By the time of Run IIb, there is expected to be a consolidation in ORACLE support by the Laboratory that will not include the existing DØ Online database Compaq / Tru64 Unix platform. These platforms will thus need to be replaced.

1.5 Interaction with the Computing Division

The Run IIa Online system was developed through an active partnership with the Computing Division's Online and Database Systems (CD/ODS) group. It is essential that this relationship be maintained during the transition to the Run IIb system. While the level of effort expended by CD/ODS personnel has already decreased relative to what it was during the height of the software development phase of the Run IIa Online system, the continued participation of this group will be needed to maintain the system and to migrate the existing software to new platforms as these are acquired. Computing Division assistance and expertise will be particularly critical in the area of database support since the Oracle consultant who led the design of the current system is not expected to be involved in maintaining the system. The continued involvement of the CD in the Online effort, which will presumably be described in a future MOU, will be left

mostly implicit in later sections of this document, but will nevertheless continue to be crucial to the success of the effort.

1.6 Comments on Run IIa configuration

For Run IIa many of the Online functions have been combined on a single cluster of large servers. The Data Logger, Database, and File Server operations are all performed on one or more of a set of three Compaq AlphaServers configured as a TruCluster. The cluster configuration allows each of the three nodes to share disk resources, which include the event data buffer disks, database disks, and general user disks. The cluster acts as an NFS server for the control room and monitoring nodes, and as the NIS master for Online accounts. High availability and reliability for these functions are provided by the cluster configuration, allowing one of the three nodes to go down with the remaining nodes assuming the critical functions.

The Run IIa cluster includes a dual-processor AlphaServer 4000 (purchased in 1997), a second dual-processor AlphaServer 4000 (1998), and a quad-processor AlphaServer GS80 (2000). Disk storage shared among the cluster members includes 1.1 TB in a fibre channel RAID array, 2.8 TB in fibre channel JBOD, and 0.6 TB in a shared SCSI RAID array. Backups are performed on a single DLT4000 tape drive on each cluster member.

The remaining Online Host, Control Room, and Monitoring nodes are single or dual processor Linux systems. These nodes are mostly interchangeable, with only the Control Room systems being slightly unusual with their configuration requiring multiple graphics cards and monitors. Other than system and scratch disks, these systems get all of their storage resources from the AlphaServer cluster.

The cluster configuration has proved both reliable and efficient. However, the I/O performance of the AlphaServer components, because of both age and architecture, is not optimal. The existing Run IIa components are maximally utilized in order to provide the Run IIa target event data rate of 50 Hz. Concentrating a large number of functions in a small number (3) of machines leads to potential congestion. The addition of resources to this specialized configuration is possible but costly.

The Run IIb architecture addresses the performance and expandability / flexibility issues by replacing the central cluster with a larger number of dedicated function Linux systems. There is a philosophy of one machine per function, with redundant systems where necessary. Where possible, functions are spread across multiple parallel machines. The details of the architecture are described in following sections.

2 DAQ/Online Computing Upgrades for Run IIb

A description of planned upgrades follows for each component noted in the Introduction. The philosophy and architecture of the Online system will not change, but components will be updated. Note that most changes are best achieved by a continuous, staged approach, while others involve large systems that will need to be replaced as units.

2.1 Operational Parameters

Table 1 summarizes the Run 2 parameters relevant to the Level 3 and Online systems. The impact of each value upon the system configuration will be discussed in the following sections.

Table 1. Run 2 Level 3 and Online parameters

Parameter	*Run IIa*	*Run IIb*
Level 3 farm nodes (dual processors)	112	160
Average event size	~ 250 Kbytes	~ 300 Kbytes
Level 3 input rate	1000 Hz	1000 Hz
Peak Level 3 accept rate	50 Hz	100 Hz
Peak logging rate	12.5 Mbytes/sec	30 Mbytes/sec
Detector duty factor	> 99%	> 99%
Accelerator duty factor	~ 75%	~ 75%
Online system availability	> 99%	> 99%
Local data buffer	48 hours	48 hours
Local data buffer	~ 2 Tbytes	~ 4 Tbytes

2.1.1 Level 3 parameters

The required number of Level 3 farm nodes is a function of the event rate into Level 3 and the required processing time for filtering. These are highly tunable numbers. The input rate can be adjusted with thresholds in Level 1 and Level 2. The processing time depends upon the choice of software filters. If a required Level 3 rejection rate is not possible in the available time, then thresholds can be adjusted at Level 3. The choice in Table 1 of 160 Level 3 farm nodes for Run IIb results from the best estimate of the required analysis time with the target input rate.

2.1.2 Event rate parameters

The rate at which events are logged determines the required capabilities of the Host Online system. The rate is bounded by the maximum trigger rejection ratio and the maximum Offline storage, reconstruction, and analysis capacities.

The Online system must be designed to cope with the largest rate otherwise allowed. This limitation is from the Offline computing systems, which see a Run IIb rate of 100 Hz with a 75% overall duty factor as the *maximum* allowable. The average output design rate as mentioned in the trigger overview is 50 Hz.

2.1.3 System availability

The DØ experience in Run1 was that an uptime of > 99% for the computing systems is achievable. In Run2 there is a similar goal, such that lost beam time from computing problems is small compared to other sources of detector down time (the Detector duty factor). Since the DØ event data is transferred to the Feynman Computing Center for logging, problems in this step can potentially contribute to system down time. To decouple possible problems with the tape robot and database systems, which are principally Offline systems with competing priorities, we require a local disk buffer to retain data for a period of 48 hours. This period should be sufficiently long to recover from any Offline system disruptions.

2.2 Online Network

2.2.1 Description

The backbone of the DØ Online computing system is the switched Ethernet network through which all components are interconnected. The Run IIa network is based on a pair of Cisco 6509 switches, one for the event data path from the SBCs in the VME readout crates to the Level 3 filter farm nodes, and the other which services the Level 3 event data output as well as general network traffic. Each switch is composed of a chassis with an interconnecting backplane and various modules that supply ports for attaching the DAQ and Online nodes. The total capacity of each switch is determined both by the chassis version and the number and versions of the component modules. The DAQ switch has a fabric-enabled backplane capable of a total throughput of 128 Gbps. The Online switch has a backplane capable of a total throughput of 64 Gbps.

The Cisco 6509 for the DAQ network can currently support 16 Gb fiber connections from the SBCs (via Cisco 2948G switches in the counting house) and 96 100Mb connections to Level 3 farm and support nodes. The existing 48 farm nodes, a pending 32-node farm addition, and a handful of support nodes will fully utilize available ports. When the number of farm nodes is increased beyond 80, additional 48-port 100Mb blades will need to be purchased for the switch.

The Online Cisco 6509 switch currently directly supports over 100 nodes, including 48 Level 3 nodes. There are approximately 48 available 100Mb ports, 32 of which will be used immediately for a Level 3 farm addition. More 100Mb blades will be needed to expand beyond the current level. An increase in the number of high bandwidth host system nodes will require more gigabit ports beyond the current fully-utilized 10 ports, necessitating the addition of a Gb capable blade.

2.2.2 Run IIb Upgrade Plan

To support a total of 160 Level 3 nodes, two additional 100baseT modules are required on each switch. To support the expanded Host Data Logging systems (page 458), one additional 1000baseT module is required in the Online switch. If a 1000baseT module is not available, existing 3COM 4900 switches (with 24 available 1000baseT ports) can be used with 1000baseSx uplinks. In this case an additional 1000baseSx module would be needed.

2.3 Level 3 Linux Filter Farm

2.3.1 Description

The Level 3 trigger consists of two principle elements: a high speed data acquisition system that provides readout of the entire detector at rates expected to exceed 1 kHz, and a processor farm that utilizes software filters written to select events that will be permanently recorded. Since the required Run IIb data acquisition bandwidth is expected to be made available once the Run IIa Level 3 hardware is fully commissioned, the most likely need for Level 3 upgrades will be to provide increased processing power in the farm.

The Run IIb Level 1 and Level 2 triggers will have increased selectivity and there will be a matching increase in selectivity of the Level 3 filter. This requires the use of more complex algorithms that necessitate the need for faster processors in the Level 3 nodes. Historically, DØ has equipped the Level 3 farm with the fastest processors on the market within this chosen processor family and this approach must be continued. At the time of the Run IIb upgrade, Moore's law would lead us to expect a four-fold increase in processing speed over what is currently available. Thus, a significant increase in Level 3 processing power will be obtained by enhancing the Run IIa Level 3 processors with the latest technology available in 2004.

There are currently 48 dual-processor Linux nodes in the Run IIa Level 3 filter farm. An additional 32 are to be installed by late 2002.

2.3.2 Level 3 Processing needs for Run IIb

The Level 3 event processing time has been measured in Run IIa running at peak luminosities of $\sim 2 \times 10^{31}$ $cm^{-2}s^{-1}$. On a 1 GHz PIII processor (~40 Si95), the filtering code with global tracking takes approximately 300 msec/event before any emphasis on time optimization. This maps into an approximate time of 250 msec/event for events with only one interaction. In the language of computing benchmarks, this translates into about 10 Si95-secs/event.

The amount of processing needed per event depends on the complexity of the event, which in turn depends on luminosity and bunch crossing time. Monte Carlo studies coupled with data analysis shows that the reconstruction time rises with the number of interactions in an event. For offline code with detailed tracking, this rise is much faster than linear mostly due to combinatorics in the global tracking. For the Online Level 3 code, we only need a limited reconstruction, depending on the lower level trigger results, to filter the event.

This reconstruction depends more on detector occupancy than event complexity, thus we assume a linear rise in processing requirements per number of interactions. With this assumption then the amount of processing power needed for 5 background interactions/bunch crossing and 1kHz Level 3 input rate leads to a predicted CPU requirement of 60k Si95.

The above argument gives us an idea of the approximate processing power needed for the Level 3 filter, however a precise processing requirement is difficult to predict. For example, with the ability to trigger on secondary vertices, the Level 3 filter code will additionally perform secondary vertex analysis. With some trigger rejection moving into Level 1 and Level 2, Level 3 may also need to work harder to get the desired rejection ratio.

2.3.3 Run IIb Upgrade Plan.

The Level 3 farm is very similar to the processing farms used in the Offline event reconstruction and analysis. A standard configuration is a 2U rack-mounted system with dual processors and 1 Gbyte of memory. The Level 3 nodes benefit from dual 100 Mb network connections to segregate input and output event data networks. The evaluation and purchase steps for Level 3 farm nodes will closely follow Offline farm activities.

Large farm purchases at Fermilab over the last 10 years have shown that processing power for commodity (desktop-like) computing has doubled every 1.5 years (Moore's law), and that prices/box remain roughly the same. This scenario is used universally inside and outside Fermilab to predict computing prices in the next 5-10 years.

Table 2 provides a scenario for the staged acquisition of farm nodes. Beyond the existing node purchases of FY01 and FY02, Moore's law is applied to predict the processing power increase relative to a dual 1 GHz machine. With an acquisition of 32 nodes per year through 2004, an eventual retirement of older nodes, and a bulk purchase of 96 nodes immediately prior to Run IIb, a farm of 160 dual-processor nodes (320 processors) with 3-4 times the Run IIa performance is expected by 2006.

Table 2. Level 3 processor acquisition schedule

Year	FY01	FY02	FY03	FY04	FY05
Added nodes	48	32	32	32	96
Node type	Dual 1 GHz PIII	Dual 1.67 GHz AMD	2003 model	2004 model	2005 model
Clock factor	1.00	1.67	2.67	4.28	6.84
Equivalent Dual 1 GHz PIII	48	53	85	136	656
Retired nodes				48	32
Node type				Dual 1 GHz PIII	Dual 1.67 GHz AMD
Clock factor				1.00	1.67
Equivalent Dual 1 GHz PIII				48	53
Net number of nodes/ports	48	80	112	96	160
Net equivalent Dual 1 GHz PIII	48	101	186	274	877

In this scenario, the total installed Level 3 computing for Run IIb is 70k Si95 compared to a need of 60k Si95. This plan is well matched to Level 3 needs given the caveats mentioned above.

2.4 Host Online Systems

2.4.1 Description

The Online Host nodes perform the functions of the Collector, Data Logger, Data Distributor, and the processes for shipping event data to the Feynman Computing Center (FCC) for recording on tape. The Collector (one or more processes acting in parallel) receives event data from the Level 3 filter farm nodes and routes it to Data Logger and Data Distributor processes. The Data Logger (one or more processes acting in parallel) writes the event data to the buffer disks. The Data Distributor maintains an event pool for monitoring applications and transmits the events to the various monitoring nodes. The DLSAM and DLCAT processes read the event data and metadata from the buffer disks and ship the information over the network to the FCC.

The Host system must be able to absorb over the network the maximum recorded data rate of 30 MB/s. This rate must be sustained as events are logged to the buffer disk, read back from the buffer disk, and shipped over the network to the FCC. Additionally, a fraction of the event stream must be routed to the event monitoring tasks. There is additional network traffic internal to the Host system as events are routed among parallel tasks.

The Host system is an integral and required component of the experiment. As such, it must be a highly available system. The target availability is greater than 99%.

2.4.2 Run IIa Host Online system

The current DØ Online Host system is centered upon three Compaq / Digital AlphaServers in a cluster configuration. Two of the machines are AlphaServer 4000s (purchased in 1997 and 1998) and the third is an AlphaServer GS80 (purchased in 2000). These machines mount disks in the form of two RAID arrays, ~600 GB in a Compaq/Digital HSZ50 unit and ~1.1 TB in a Compaq/Digital HSG80 unit, and an additional 2.8 TB in Fibre Channel JBOD disk. This cluster supports data logging, the ORACLE databases, and general file serving for the remainder of the Online system.

The long-term maintenance of these systems is a serious concern. While they can be expected to still be operational in the Run IIb era, the high availability required for critical system components may be compromised by the inability to obtain the necessary maintenance support. Maintenance costs for these systems, particularly 7x24 coverage, will increase with age. By the time of Run IIb, maintenance costs are likely to exceed replacement costs.

These systems currently run Compaq Tru64 UNIX, previously known as Digital UNIX, or Digital OSF1. With the merger of Compaq and Hewlett Packard, long-term support for this operating system may be problematic.

2.4.3 Run IIb Upgrade plan

All applications developed for the data acquisition system that currently run on the Host systems were written with portability in mind. In particular, all will

work under Linux. The proposed upgrade to the Host systems is therefore to replace them with Linux servers. Since the existing Host system provides Data Logging, Database support, and File Serving functions, each of these needs must be accommodated by the replacement system. These requirements will be addressed individually in this and following sections.

The Data Logging system must, with high (> 99%) availability, be capable of absorbing data from the Level 3 filter systems, distributing it to logging and monitoring applications, spooling it to disk, reading it from disk, and dispatching it to tape-writing nodes in the FCC. The maximum input data rate is 30 Mbytes/sec. A local disk buffer of ~4 Tbytes is required to retain event data if the Offline tape robots are unavailable. The event data on the buffer disks is normally read and transferred to the FCC at the 30 Mbytes/sec input rate, but a 2x higher rate is necessary to both take sustained new data and unload stored event data following any outage. The high availability requirement, satisfied in the current system by using a cluster of three machines, precludes the use of a single machine. Currently the cluster members share the disk buffers, but this is not a strict requirement.

The proposed upgrade solution is for a set (two or three) of Linux servers (dual or quad processors) to act as the new Data Logging nodes. The data acquisition applications can run in parallel to distribute the load at full bandwidth, but a single node should be capable of handling nearly the entire bandwidth for running under special conditions. Each system will require gigabit connectivity to the Online switch, thereby raising the number of gigabit ports required.

Some R&D effort is needed to test such a configuration. The possibility of clustering the Linux nodes and the possibility of sharing the disk storage will be examined. The purchase of the complete Data Logging system can be staged, as not all members need to be identical (as noted above, the current Host system was purchased in three increments). It is expected that the R&D unit to be purchased could suffice as the third server in a three-server configuration.

The expected configuration consists of three high-end servers with a SAN JBOD disk array. A typical server would be configured with four CPUs, 4 Gbytes of memory, a local RAID controller for system disks, gigabit network cards, and available 64-bit PCI slots. The buffer disk array will be built from commodity disks in a JBOD Fibre Channel crate, as currently employed for the Run IIa system.

Other components of the Host Online system – the Collector and Distributor – will be housed on Run IIa Linux systems which will remain sufficiently capable for these functions in Run IIb.

2.5 Control Room Systems

2.5.1 Description

The current DØ Control Room system is composed of 12 Linux nodes (single and dual processor) that manage 27 monitors. These systems range in age from one to five years. Many of the monitors are already showing the effects of age.

It is expected that we should replace some fraction of the Control Room nodes and monitors each year.

2.5.2 Run IIb Upgrade plan

The Control Room systems will be gradually upgraded and replaced throughout the lifetime of the DØ experiment. Assuming a 5-year viable lifetime for these components, we should expect to replace 20% each year. This implies that 2 to 3 systems supporting 5 to 6 monitors will need to be replaced annually. This will be done with purchases of moderate-performance (dual CPU, multiple graphics cards) graphics systems or reuse of systems otherwise made available.

2.6 Data Monitoring Systems

2.6.1 Description

Real-time monitoring of event data is accomplished by a scheme in which representative events are replicated and distributed to monitoring nodes as they are acquired. The monitoring ranges from examination of low-level quantities such as hit and pulse height distributions to complete event reconstruction. In the latter case, the environment and the code are similar to that of the Offline reconstruction farms. There are one or more monitoring applications for each detector subsystem, and for the trigger, luminosity, and global reconstruction tasks.

The rate at which the monitoring tasks can process events, as well as the complexity of monitoring, is limited by the processing capabilities of the monitoring nodes. The Control Room systems and several rack-mounted Linux nodes currently share this load. Much can be gained by upgrading the experiment's monitoring capability. As more sophisticated analysis software becomes available, these improved codes can be run in the Online environment to provide immediate feedback on data quality.

2.6.2 Run IIb Upgrade plan

The monitoring nodes, rack mounted Linux systems, will need to be continually updated. Such upgrades can occur gradually. As with the Control Room nodes, a useful lifetime of 5 years implies that 20% of the monitoring systems should be upgraded each year. A typical monitoring node configuration is a 2U rack-mounted dual-processor system with 1 GB of memory. These systems are very similar to the Level 3 or Offline farm nodes, and purchases will be made in conjunction with these larger acquisitions. Four or five nodes will be upgraded per year.

2.7 Database Servers

2.7.1 Description

The ORACLE databases currently run on the AlphaServer cluster, with the database files residing on the attached RAID arrays. As mentioned above, long-term support for this hardware is questionable. Additionally, ORACLE database and application support from the Computing Division no longer includes the Tru64 UNIX platform.

The principal requirement for the database server is high availability (> 99%). Support needs include maintaining the hardware, the operating system, and the application software (ORACLE). User application development also benefits from having independent production and development database instances.

2.7.2 Run IIb Upgrade plan

The planned replacement of the database servers is by two redundant Linux systems with common access to RAID disk arrays. The Computing Division supports this system. The purchase of these systems is best staged over two years, with early purchase of the development machine and later purchase of the production machine. The performance required of the database machines is not expected to be extremely demanding. A mid-range server system should be sufficient.

The RAID array for the database will likely be shared with the File Server systems, as neither application has particularly demanding disk I/O performance. The production database instance is expected to require no more than 400 Gbytes, and 100 Gbytes should suffice for the development instance (Run IIa sizes are 80 Gbytes and 35 Gbytes). An additional 200 Gbytes will be necessary as spooling space for the ORACLE backups.

The Database system will be supported with a backup system, in conjunction with the File Server systems of the next section.

2.8 File Servers

2.8.1 Description

The Host cluster currently provides general-purpose file serving. Linux nodes within the Online system access the Host file systems by NFS. Approximately 500 GB of RAID disk is currently available. Files stored include the DØ software library, Fermilab software products, DAQ configuration files, detector subsystem application data, and user home areas. Since the existing file servers are the AlphaServers, replacement is necessary, for reasons already delineated.

The requirement for the file server system is primarily one of high reliability (> 99%) of both system and disks. The needed network and disk I/O rate is moderate, with 5 Mbytes/sec being sufficient and easily achievable.

2.8.2 Run IIb Upgrade plans

The proposed solution is a pair of redundant Linux servers with common access to both RAID and JBOD disk arrays. The RAID array will be shared with the Database systems. Assuming a 1.3 TB RAID array and the previously stated database needs, then 600 GB will be available for general system use. The JBOD disk is typically provided from otherwise unused disk space on local disks of the distributed nodes, so no explicit purchase is necessary. A tape stacker system of moderate capability is needed for backups. Acquisition of these systems can be staged.

2.9 Slow Control Systems

2.9.1 Description

The Slow Controls system consists of host-level console computers that run the controls application programs and that communicate with Input/Output Controller (IOC) computers over an Ethernet LAN. The IOC's, in turn, communicate with detector hardware components (many of which have their own embedded process control computers) over two types of field bus: (a) the VME parallel backplane and (b) the MIL/STD1553B serial bus.

An IOC processor is either a Motorola 68K or PowerPC single-board computer. The operating system for these processors is VxWorks and the controls-specific software consists of EPICS (Experimental Physics and Industrial Control System) plus a number of EPICS extensions developed by the DØ controls group. The node-specific EPICS configuration, which is loaded into an IOC at boot time, is maintained in an ORACLE database from which the IOC configuration files are extracted. The processors have both Ethernet connections to the Online network and serial line connections from their console ports to network terminal servers. The console connections allow users to communicate directly with the VxWorks shell and the EPICS local user interface.

The IOC nodes perform downloading, monitoring, calibration, and other time-critical functions essential to the operation of the detector. The host-level nodes execute applications such as operator monitoring and control GUI's, the detector configuration manager, the central significant event (alarm) system, and other, less time-critical tasks.

Multiple MIL/STD1553B busses provide the communication link to all the electronic components on the platform and many of the components in the movable counting house. Controller cards that are located in VME crates containing an IOC processor drive these busses. The MIL/STD1553B bus was selected for its robustness and for the low electrical noise environment required for devices located on the detector platform.

A significant fraction of the sensors that are managed by the slow controls system are connected via a generic analog/digital interface unit called a rack monitor (RM) that communicates with an IOC via a MIL/STD1553B bus. The environmental conditions in electronics racks on the platform and in the movable counting house are monitored by a Rack Monitor Interface (RMI) unit that, in turn, is connected to a RM.

2.9.2 Run IIb Upgrade Plan

Both the Motorola 68K and PowerPC processor families used for Run IIa have limited lifetimes. By the beginning of Run IIb, repairs or replacements for the 68K processor boards will no longer be available and, by the end of the run, the same situation may also exist for the PowerPC boards. Furthermore, the functionality of the 68K processors in the Muon detector read-out crates is now severely limited by their available memory (4 Mbytes). Due to the age of these processors, memory upgrades are no longer available. Monitoring, control, and

calibration functionality would be greatly improved by a complete replacement of these aging processors.

To remedy this situation, we have developed a plan for replacing several groups of processors in a manner that addresses performance needs and requirements for spares. Table 3 notes the replacement scheme.

Table 3 Controls processor upgrade plan

Detector subsystem	# of Processors	Processor types	Replacement plan
Control and Monitoring	12	(6) 16MB PowerPC (5) 64MB PowerPC (1) 128MB PowerPC	Replace, use old processors for HV or spares
High Voltage	29	(29) 16MB PowerPC	Retain, with new and spare needs met from other replacements
Muon	~40	(40) 4MB 68K	Replace (17) in readout crates, retain as spares
Tracker readout	24	(12) 64MB PowerPC (12) 128MB PowerPC	Replace, use old processors for HV or spares
Test stands	~7	mixed low end	Use available

The replacement processors will be selected on the basis of performance, cost, and support by the EPICS community. They are expected to be very similar to, if not identical, the SBC processors used for the DAQ readout. Long term spare and maintenance support will be eased if the components are the same.

The DØ Electrical group supports the MIL/STD1553B controllers, Rack Monitors, and Rack Monitor Interfaces. They expect to produce new units sufficient to supply the needs of Run IIb for new systems and spares.

3 DAQ/Online Computing Summary

The need to update and replace DØ DAQ and Online Computing equipment is based mainly on the problems associated with the rapid aging and obsolescence of computing hardware. Maintenance costs, particularly 7x24 costs for high availability systems, rapidly approach replacement costs by systems with much greater functionality. Additionally, software support for operating systems and critical applications (ORACLE) is potentially problematic for the platforms currently in use. There is a need for higher bandwidth data logging made possible by improved Offline capabilities. There are very real benefits to be accrued from more complex trigger filters and data monitoring software. For these reasons, we plan to update and replace the Online systems.

Replacement systems, wherever possible, will be based on commodity Linux solutions. This is expected to provide the best performance at the lowest cost. The Fermilab Computing Division is expected to support Linux as a primary operating system, with full support of local products and commercial applications. We plan to follow a “one machine, one function” philosophy in organizing the structure of the Online system. In this way, less costly commodity processors can replace costly large machines.

The maintenance of the Control System will be difficult in the Run IIb era as the component parts become difficult to acquire. The plan is to replace a number of processors, using the replaced components as spares for the remainder.

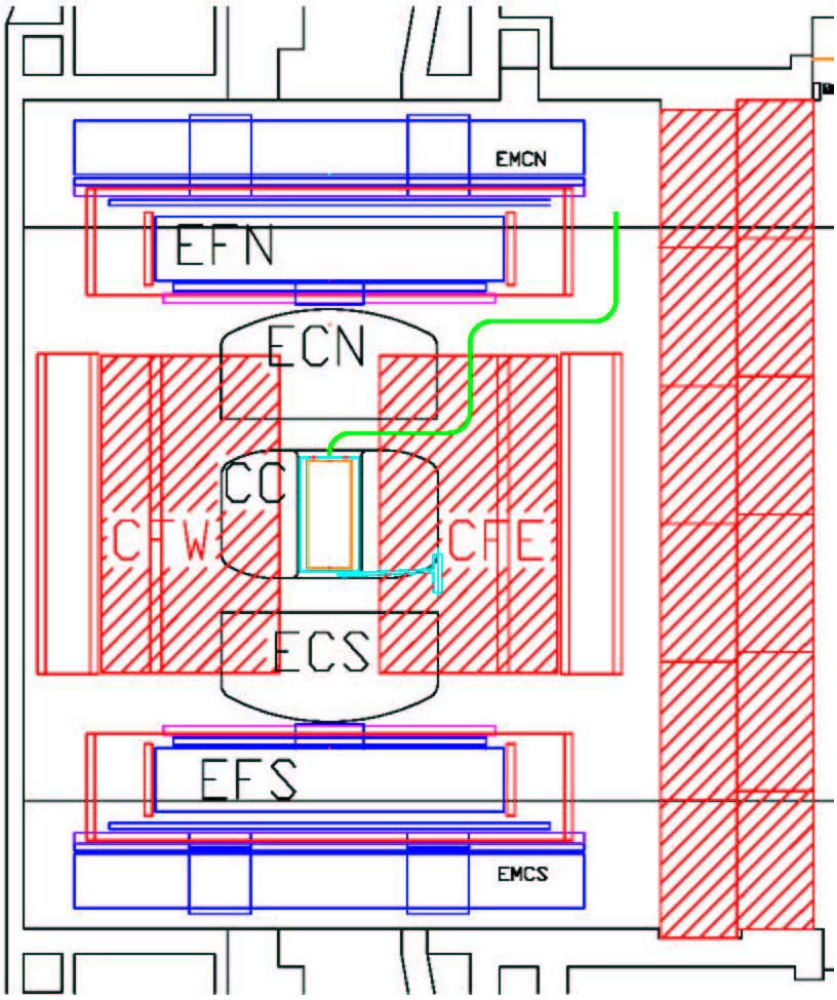

INSTALLATION

(This page intentionally left blank)

INSTALLATION CONTENTS

(This page intentionally left blank)

1 Installation Introduction

1.1 Overview

A new silicon microstrip detector will be fabricated for the DØ detector for Run IIb. Upgrades to the DØ trigger system will also be prepared for Run IIb. It is important to conduct the installation of these systems to minimize the shutdown period of the Tevatron collider between Run IIa and Run IIb.

Because the new silicon detector has been designed to part at z = 0, each of the two halves of the system are sufficiently compact to permit them to be installed without rolling the DØ detector out of the collision hall. This feature will save substantial time and effort during the shutdown. Also, the Run IIa silicon signal cable plant and high voltage system will be reused for Run IIb, saving substantial effort and time during the installation period.

The majority of the new trigger system elements will be installed in the DØ Moveable Counting House (MCH), so they do not require substantial access to the detector during the shutdown.

The DAQ/Online computing system will also receive substantial upgrades for Run IIb. Because these upgrades consist largely of enhanced-capacity single components that will be installed over time without significant change to the architecture of the Run IIa system, their installation can be conveniently managed along with general DAQ/Online system support and maintenance. Thus their installation will not be discussed herein.

Finally, certain of the infrastructure needs of both the silicon and trigger systems can be prepared in advance so installation and commissioning time are minimized.

1.2 Scope

The Run IIb Installation task will concern itself with the installation and technical commissioning of the elements of the DØ detector that have been upgraded for Run IIb:

- Silicon Microstrip Detector,
- Level 1 Calorimeter Trigger,
- Level 1 Calorimeter Track Matching,
- Level 1 Central Track Trigger,
- Level 2 β Processors,
- Silicon Track Trigger,
- Plus the associated detector infrastructure for each of these elements.

For the Installation task, “technical commissioning” means that after an upgrade element is installed, it will be powered up and performance-tested until it is deemed “ready for beam”, i.e. the DØ detector can be closed and the collision hall secured because daily access to the element in question is no longer required. For example, for the silicon detector, technical commissioning includes demonstrating that all components power up safely to design voltage and current

and that all actively cooled components stabilize at desired temperatures. Technical commissioning also includes demonstrating that all crates can be downloaded, that all electronic channels are properly mapped to physics channels, that noise, pedestals, (even cosmic rays if appropriate) are understood, that simulations using stored patterns are processed properly by the system, and that the silicon software is sufficiently mature to operate the system as required. Technical commissioning ends when the system is deemed by its commissioners to be fully ready for beam.

1.3 Schedule

Given the date of completion of the Run IIb DØ silicon detector at SiDet it is appropriate to plan the overall installation schedule so that when the tasks necessary to remove the existing Run IIa silicon system are considered, the latest access time to the detector (i.e. the shutdown date of the Tevatron) can be chosen. This date is found to be approximately 9 weeks in advance of the completion of the silicon.

The shutdown will begin with the opening of the detector and the draining of the coolant from the existing silicon, followed promptly by the uncabling of the silicon. Then the beryllium beampipe and the silicon system itself will be removed, along with the old adapter card supports ("horseshoes"). After the new horseshoes are installed, the Run IIb silicon will be delivered and installed, the cables reconnected, and the new beampipe inserted. This point will be reached approximately 10 weeks after the new silicon is delivered.

At this time, the technical commissioning of the silicon will begin because during the preceding weeks the new silicon infrastructure (interface boards, power supplies, temperature and radiation monitoring systems, cooling systems, etc.) will have been installed and hooked up. Approximately 11 weeks will be required for technical commissioning which brings the new silicon system to the state that it is ready for beam. At this point the detector will be closed and Tevatron operations may resume. The total shutdown time of the Tevatron is approximately 30 weeks.

Also at the beginning of the shutdown, the Level 1 Calorimeter trigger will be removed from the MCH and the new system installed. The other elements of the trigger upgrades will be installed, and in the ensuing weeks all new trigger systems will be brought to a state of operation such that they too are ready for beam.

The installation task is complete when the detector is closed up and ready for beam.

Note that certain of the infrastructure tasks (e.g. preparation of the racks for the L1 Cal trigger, preliminary installation of the silicon cooling system on the sidewalks outside the collision hall, testing and labeling of cables, etc.) will begin well before the shutdown of the Tevatron. These tasks are included in the installation schedule because their timely completion directly bears on the

success of the post-shutdown portions of the installation schedule. These infrastructure tasks begin approximately one year before the end of Run IIa.

1.4 Management

Four Level 3 managers (one for silicon mechanical, two for silicon electronic, and one for trigger) are presently in place who have expertise in the electrical and mechanical details of the silicon detector or are knowledgeable in matters pertaining to the trigger system. Each has Run IIa installation experience.

These managers have already begun to plan the infrastructure tasks appropriate to their positions and they will direct these activities during the last year or so of Run IIa. For example, the Level 3 silicon electronic installation manager will participate in the design of the new silicon low voltage system and cables so that installation planning for these elements is detailed and credible. As soon as all the new cable systems are defined and understood, the procured cables will be tested, labeled, bundled if appropriate, and held in full readiness for installation when the shutdown begins. Effort is included in the Installation schedule that provides for a “cable czar” who will be responsible for all aspects of cable preparation and installation. That person will provide continuity for this substantial responsibility and with a Level 3 silicon electronic manager will direct the installation of the cables. Similar considerations pertain to the silicon cooling systems, and the adapter card supports.

The Level 3 managers also have begun to plan the primary installation activities (i.e. silicon and trigger systems), and they will direct these tasks during the shutdown. During Run IIa the Level 3 managers will meet periodically with the principals (engineers and physicists) of the silicon and trigger tasks so that full understanding of the state of readiness of these systems is monitored and understood. For example, the redesign of the high voltage system will be tracked so that installation planning is relevant, detailed and credible. Likewise, the planning for the redressing of the L1 Calorimeter trigger cable plant will be thorough and fully organized so when the shutdown begins the work can proceed with minimal chance of significant upset to the schedule.

2 Installation Specifics

2.1 Run IIa Silicon Removal

Immediately after the opening of the detector, the “low mass” cables from the existing silicon detector will be disconnected from the adapter cards on the horseshoes and the beryllium beampipe will be removed. This beampipe will be withdrawn out through the beampipe bore of an end calorimeter. The 80-conductor “high mass” cables from the adapter cards will be removed from the horseshoes and carefully draped aside on the face of the central calorimeter (CC) for reuse. The H-disks of the old silicon tracker will then be removed from the Central Fiber Tracker (CFT), and both the north (N) and south (S) halves of the silicon tracker removed in turn. To ensure no damage to the CFT, the installation fixturing used to install the silicon for RunIIa will be reused to carefully remove the silicon detector halves. While the N silicon is withdrawn, the S horseshoe will be demounted, and vice-versa.

Also as soon as the detector is opened, one crate’s worth of interface boards will be removed from one IB crate beside the CC. These boards will be express-shipped to KSU for reworking for Run IIb. Because it is estimated to take three weeks to modify and test these boards, it is important to get the KSU facility “up and running” promptly. Helpfully, one crate’s worth of boards will already have been modified at KSU for use at SiDet much earlier. The balance of the eight IB crates will be removed and emptied over a two-week period and the boards shipped to KSU to be folded into the modify/test schedule there. As each IB crate is removed the 80-conductor cables that fed it will be tested and certified for reuse in RunIIb. The low voltage (LV) power supplies and ancillary components above the IB crates will be removed at this time as well.

2.2 Run IIb Silicon Installation

The new horseshoes will be installed on CC and the precision mounts installed in the CFT that support the new silicon. As soon as the latter are verified for proper alignment, the silicon will be installed in CFT. As was done during Run IIa, the two halves will be installed separately using a trolley and table. Novel for Run IIb however, the z=0 joint will be made up by sliding both halves towards an EC to expose the fasteners at the z=0 joint. The joined halves will then be moved back in z to the center of the CFT and the end mounts made up and checked for alignment. The cooling and dry gas systems will be connected and put into operation, and the milestone “Silicon installed in CFT” will have been achieved.

While work inside the detector has proceeded, the new low voltage system and augmented high voltage (HV) systems have been installed in the MCH and on the detector platform and checked.

In the detector itself, after the silicon is in place, the new junction cards (which carry the new twisted-pair signal cables) will be installed in the bore of CFT and the flex-cables from the silicon connected to them. The original 80-conductor cables will be connected to the new adapter cards on the horseshoes,

and the new beryllium beam tube will be installed in the silicon. Reworked and tested interface boards will have steadily returned from KSU and will have been installed in IB crates in sequence.

2.3 Technical Commissioning of Silicon

Inasmuch as the HV and LV power supply systems will have by this point been cabled to the IB crate area and horseshoes, and the cooling system will be up and running, and the detector fully cabled, the technical commissioning of the silicon will begin, quadrant by quadrant as the IB crates are loaded . As the fraction of the detector that is commissioned grows, the online and level 3 software will be tested, and all hardware channel to physics channel mappings will be verified.

When the technical commissioning is complete, the detector is closed and the silicon system is ready for beam.

2.4 L1 Calorimeter Trigger System

Immediately at the beginning of the shutdown, the 13 Level 1 Calorimeter Trigger racks will be uncabled and removed from MCH1. The L1 Cal trigger cables underneath the floor in MCH1 will be extracted from the very limited volume where they were installed prior to the beginning of Run I nearly 15 years ago, and the bundles undone, sorted, redressed so they will reach the new trigger crates, and tested. When this very time-consuming job is completed, the new racks (prepared and tested prior to the shutdown) will be brought into MCH1 and fastened in place. The new trigger crates containing the ADF's (ADC with Digital Filtering), TAB's (Trigger Algorithm Boards), and GAB's (Global Algorithm Boards) will be reinstalled in the racks, all rack services reconnected, and the very extensive technical commissioning of L1 Cal trigger will begin.

It should be noted that critical services (chief among which is the VESDA fire detection system) for the trigger framework (TFW) will be maintained from the beginning of the work in MCH1 so that the TFW remains available throughout the shutdown (except for perhaps a day or two at the beginning). This means that work on the DAQ, as well as commissioning activities for the silicon and the trigger, can proceed when otherwise ready.

2.5 L1 Cal Track Match

After a modest period of installation for the two new VME crates with new trigger components (SLDB's – Serial Link Daughter Boards), MTCxx's and MTCM's (L1 Track and Cal Match), and MTFB (Muon Trigger Flavor Boards), the extensive system tests of the new Cal Track Match system can begin as soon as one-eight of the L1 Cal Trig system is up and running.

2.6 L1 Central Track Trigger

After a short period of installation of the new DFEA's (Digital Front End Axial) daughter boards in existing crates, technical commissioning of this system can begin, since the TFW will be available soon after the shutdown begins and the

balance of the DAQ is not expected to experience significant downtime as it too is upgraded

2.7 L2 Beta Processors

The new beta SBC's (single board computers) will be direct replacements of existing SBC's in the L2 Trigger crates, so installation of these elements will be straightforward. Commissioning of these new processors will be completed early in the shutdown.

2.8 Silicon Track Trigger Upgrade

An additional 20 VME motherboards, with a variety of daughterboards and special-function modules mostly from Run IIa, plus repeaters, and additional splitters and fiber cables, will be installed in MCH2. Due to the simplicity of this work, it will be completed promptly, but technical commissioning of the STT upgrade cannot finish until the L1 Cal Track Match is substantially commissioned and the silicon is also substantially commissioned. These latter two constraints mean that the STT will not be ready for beam until late in the shutdown.

3 Installation Resources and Schedule

3.1 Resources

Installation effort by technicians and physicists dominates the cost of installation. With modest leveling a peak of 17.5 FTE technicians and a peak of 19.1 FTE physicists are required during the shutdown. Over the 8 months the average number of technicians is 9.5 FTE's and the average number of physicists is 13.7 FTE's. The peak for electrical engineers is 3.2 FTE's and for mechanical engineers, 2.3 FTE's.

3.2 Schedule

The schedule (Table 1) developed for the installation period necessarily keys to the completion date of the silicon at SiDet. Additional trigger element completion dates are also inherited, but they are not found to constrain the installation schedule. These milestones in Table 1 are shown in italic to distinguish them from those internal to the installation itself, and red text guides the eye to key milestones.

Table 1. Installation milestones.

Milestone	Date
L1 Cal/Track Match Production and Testing Completed	*6/16/04*
STT Hardware production complete	*12/8/04*
Level 2 Beta Production Complete	*1/28/05*
L1 Cal TAB/GAB Production And Testing Complete	*2/1/05*
L1 Central Track Trigger DFEA Production And Testing Complete	*3/9/05*
Shutdown for Installation Begins	5/25/05
Detector Open and Ready for Access	6/1/05
Run IIa Silicon H-Disks Removed	6/16/05
Run IIa North and South Silicon Detectors Removed	6/30/05
STT Hardware Installed	7/7/05
Level 2 Upgrade Beta Installed & Commissioned	7/11/05
Run IIa North and South Cable Horseshoes Installed on EC Faces	7/15/05
Silicon Ready To Move To DAB	*7/22/05*
Silicon Infrastructure Prepared	8/1/05
Level 1 Tracking Installation Complete	8/18/05
Detector Installed In Fiber Tracker	8/22/05
RunIIb Silicon High Voltage System Installed	9/1/05
RunIIb Silicon Low Voltage System Installed	9/9/05
Ready to begin Technical Commissioning of Silicon	9/28/05
Silicon 80-conductor and Clock Cables Installed	10/4/05
Run IIb Beam Tube Installed	10/18/05
Level 1 Calorimeter Trigger Installed	12/20/05
Level 1 Calorimeter Track Matching Installation Complete	12/20/05
Silicon Track Trigger Operational	12/20/05
Trigger Upgrade Ready for Beam	12/20/05
Silicon System Ready for Beam	12/22/05
Run IIb Detector Ready for Beam	12/22/05

4 Installation Summary

The technical details of the installation of the RunIIb silicon will not vary greatly from those of the Run IIa silicon. Hindsight is a good teacher, and it is believed that with careful advance planning and thorough preparation of infrastructure (cabling, etc.), plus the recycling of much of the Run IIa cable plant, the schedule for the Run IIb installation can be significantly shorter than that experienced for Run IIa, when cabling alone required five months elapsed time, working two shifts per day. Major simplifications from the Run IIa experience include the fact that no operation of the Tevatron is expected during the installation period as happened for Run IIa, that the high-mass cables will not be replaced, and that all electronics will be on hand and installed before cabling begins.

As the milestone table shows, the infrastructure tasks will be completed before any of the elements they provide are required, and all of the trigger upgrade electronics will be in hand before Run IIa ends. The same engineering team that prepared the silicon will be available to guide its installation, and many of the physicists who participated in the commissioning of the Run IIa silicon are expected to participate in the commissioning of the Run IIb silicon.

(This page intentionally left blank)

DØ Run IIb Upgrade Technical Design Report

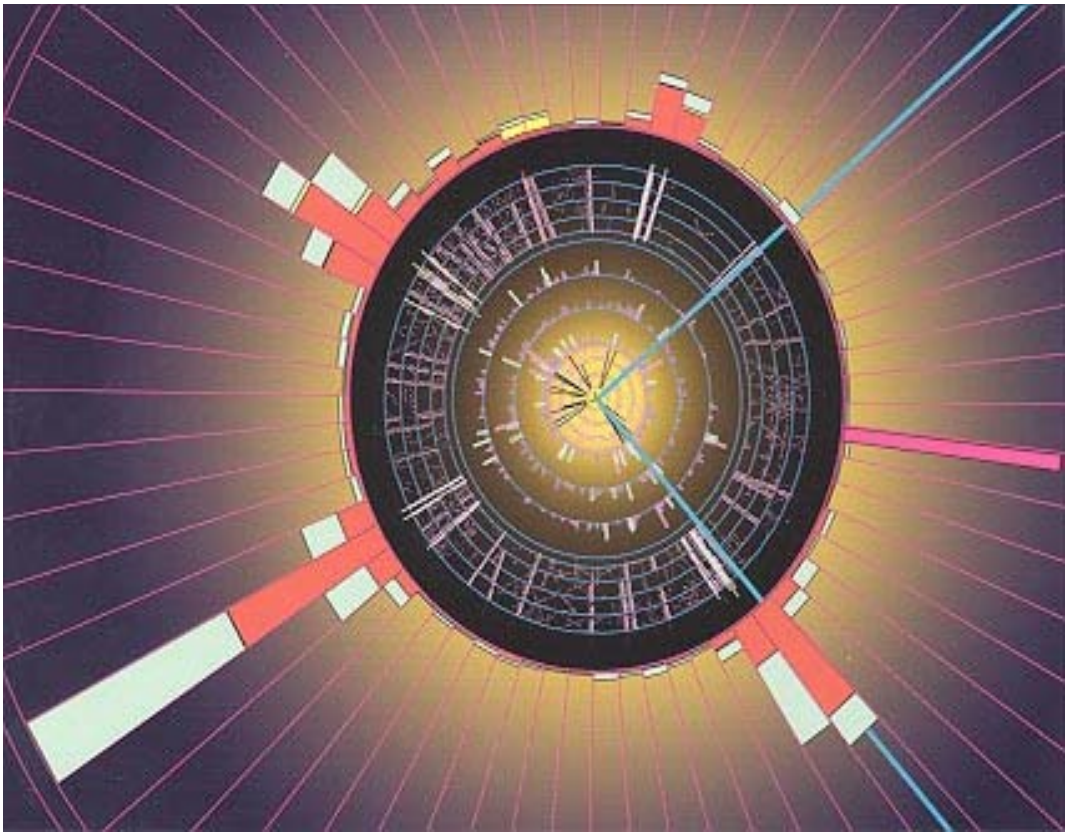

SUMMARY

(This page intentionally left blank)

SUMMARY

Clear and compelling physics goals, as outlined in the first part of this report, are the driving force behind the DØ Run IIb upgrade. The Fermilab Tevatron collider provides a unique opportunity during the next several years to make incisive searches for the Higgs, SUSY, and a wide variety of new phenomena beyond the Standard Model. These exciting prospects for a major discovery are complemented by the ability to make sensitive studies of top quark properties, precision electroweak measurements, and a wide range of investigations of Standard Model predictions.

Success in these goals can only be achieved by upgrading the DØ detector to meet the challenges posed by extended running at high luminosity. We present in this report our studies demonstrating the need for upgrades to the silicon detector, the trigger, and the DAQ/online computing systems. These studies establish the performance of the proposed upgrades, evaluated both in terms of technical performance and in the metric of the Standard Model Higgs search.

The bulk of this report deals with the detailed technical design of the proposed upgrades. Extensive engineering studies have been performed, and detailed designs for all major components of the Run IIb upgrade are described. An aggressive prototyping phase is underway to verify the feasibility and performance of the design. Long leadtime elements, such as development of the SVX4 chip, have been given special attention. Finally, extensive simulations have verified that the design meets the performance requirements set by the Run IIb physics program.

(This page intentionally left blank)